

---

Research paper

# Malware in the future? Forecasting of analyst detection of cyber events

Jonathan Z. Bakdash,<sup>1,2,\*</sup> Steve Hutchinson,<sup>3</sup> Erin G. Zaroukian,<sup>4</sup>  
Laura R. Marusich,<sup>5</sup> Saravanan Thirumuruganathan,<sup>6</sup>  
Charmaine Sample,<sup>3</sup> Blaine Hoffman,<sup>4</sup> and Gautam Das<sup>7</sup>

<sup>1</sup>US Army Research Laboratory South at the University of Texas at Dallas, Richardson, TX, USA, <sup>2</sup>Department of Psychology, Counseling, and Special Education, Texas A&M Commerce, Commerce, TX, USA, <sup>3</sup>Computational and Information Sciences Directorate, ICF for the US Army Research Laboratory, Adelphi, MD, USA, <sup>4</sup>Human Research and Engineering Directorate, US Army Research Laboratory, Aberdeen Proving Ground, MD, USA, <sup>5</sup>US Army Research Laboratory South at the University of Texas at Arlington, Arlington, TX, USA, <sup>6</sup>Qatar Computing and Research Institute, Qatar Foundation, Doha, Qatar, and <sup>7</sup>Computer Science and Engineering Department, University of Texas at Arlington, Arlington, TX, USA

\*Corresponding address: E-mail: jonathan.z.bakdash.civ@mail.mil

Received 1 November 2017; revised 8 June 2018; accepted 2 October 2018

## Abstract

Cyberattacks endanger physical, economic, social, and political security. There have been extensive efforts in government, academia, and industry to anticipate, forecast, and mitigate such cyberattacks. A common approach is time-series forecasting of cyberattacks based on data from network telescopes, honeypots, and automated intrusion detection/prevention systems. This research has uncovered key insights such as systematicity in cyberattacks. Here, we propose an alternate perspective of this problem by performing forecasting of attacks that are “analyst-detected” and “-verified” occurrences of malware. We call these instances of malware cyber event data. Specifically, our dataset was analyst-detected incidents from a large operational Computer Security Service Provider (CSSP) for the US Department of Defense, which rarely relies only on automated systems. Our data set consists of weekly counts of cyber events over approximately 7 years. This curated dataset has characteristics that distinguish it from most datasets used in prior research on cyberattacks. Since all cyber events were validated by analysts, our dataset is unlikely to have false positives which are often endemic in other sources of data. Further, the higher-quality data could be used for a number of important tasks for CSSPs such as resource allocation, estimation of security resources, and the development of effective risk-management strategies. To quantify bursts, we used a Markov model of state transitions. For forecasting, we used a Bayesian State Space Model and found that events one week ahead could be predicted with reasonable accuracy, with the exception of bursts. Our findings of systematicity in analyst-detected cyberattacks are consistent with previous work using cyberattack data from other sources. The advanced information provided by a forecast may help with threat awareness by providing a probable value and range for future cyber events one week ahead, similar to a weather forecast. Other potential applications for cyber event forecasting include proactive allocation of resources and capabilities for cyber defense (e.g., analyst staffing and sensor configuration) in CSSPs. Enhanced threat awareness may improve cybersecurity by helping to optimize human and technical capabilities for cyber defense.

**Key words:** cybersecurity; forecasting; prediction; cyberattack; malware; computer security service provider

## Introduction

Internet infrastructure plays a crucial role in a number of daily activities. The pervasive nature of cyber systems ensures far-reaching consequences of cyberattacks. Cyberattacks threaten physical, economic, social, and political security. The effects of cyberattacks can disrupt, deny, and even disable the operation of critical infrastructure including power grids, communication networks, hospitals, financial institutions, and defense and military systems. To protect its critical infrastructure, the US Department of Defense (DoD) has identified cyberspace (information networks for computers, communication, and other systems) as a key operational environment for the military, one that is interdependent with the physical (air, land, maritime, and space) environment [1]. A key component of the DoD's strategy and implementation plans for protecting cyberspace is enhancing threat awareness for Computer Security Service Providers<sup>1</sup> [CSSPs] [2–3]. Analysts in DoD CSSPs protect DoD and DoD affiliated computers and networks by finding, analyzing, remediating, and documenting cyberattacks.

To improve threat awareness for CSSPs, we investigate whether intrinsic, predictable patterns exist among “analyst-detected” and “-verified” occurrences of malware, referred to here as cyber events. This research is unique because the dataset comprises over 7 years of cyber events from an operational DoD CSSP that rarely relies only on automated systems (e.g., anti-virus software, firewalls, intrusion detection systems (IDS), and intrusion protection systems which are typically signature based) to detect and verify attacks.

In contrast, nearly all prior research on modeling cyberattacks [4–10] lacked analyst detection and verification of computer security incidents with the exceptions of [11, 12]. In these two exceptions, security incidents were verified by system administrators at a large university [11] or verified by analysts at a CSSP [12]. Thus in most earlier research, the sources for cyberattacks were processed data from network telescopes and honeypots [4, 6, 8–10] and alerts from automated systems on real networks [5, 7, 13]. Compared to real networks, the majority of traffic to network telescopes (passive monitoring of unrequested network traffic to unused IP addresses) and honeypots (monitored and isolated systems that are designed to appear legitimate to attackers) can be considered suspicious. However, the rates of false alarms and misses for cyberattacks depend upon how “attacks” are determined in data processing and/or the automated systems [14]. Cyberattacks inferred from automated systems generate a large volume of alerts and are considered a standard measure for attacks [15]. Yet automated systems only indicate an attack “may” have occurred because of false positives and misses [12, 16]. Many automated systems compare information gathered by sensors to patterns stored as signatures, looking for matches. A signature that is too specific may miss attacks, just as one that is not specific enough may generate false positives. Consequently, there are multiple approaches to manage the volume of alerts by filtering and combining them (e.g., flows of related alerts) [7, 13]. Other work on improving IDS detection has applied machine learning [17–18] to a widely used synthetic dataset on cyberattacks, Knowledge Discovery and Data Mining CUP (KDDCUP) 1999 [19]. The original dataset contained a high number of duplicate records,

which inflated classification accuracy and was addressed in a new version of the dataset, NSL-KDD [18]. Nonetheless, the KDDCUP 1999 dataset and derivatives of it are now quite old and are unlikely to be representative of current systems.

Despite potential limitations on the data sources in the majority of prior research, intrinsic patterns in cyberattacks have been identified. The main finding is that cyberattack frequency can be forecast over time using processed data from network telescopes and honeypots as well as automated systems [4, 6–10, 20]. Similarly, the only research using human-verified cyberattack data also found intrinsic temporal patterns [11, 12]. Other patterns have also been found, including the presence of bursts or extreme values [6, 9–10] and disproportionate exploitation of specific vulnerabilities [5]. Taken together, this work suggests that cyberattacks have a deterministic component: They are not fully stochastic (or random point processes).

The current work is novel because little is known about the systematicity for analyst detected and verified cyberattacks protecting critical infrastructure, in this case, US military networks, including organizations affiliated with the US Department of Defense. The curated analyst dataset is unlikely to suffer from false positives because cyber events are detected and verified by analysts who investigate events and connect evidence to confirm or disconfirm potential events, also see [12]. Consequently, analyst-detected and -verified cyber events provide a potentially more direct, filtered, and informative indicator of threats for CSSPs than attack data processed from network telescopes and honeypots and alerts from automated systems in isolation. However, in our data and other real-world datasets, whether with automated systems or analysts, the number of attacks that were missed is unknown. A miss could be an attack that was completely overlooked or delayed detection.

In this article, we make three crucial contributions. First, as previously described, we advocate for the forecasting of cyber incidents from analysts, as it has a number of appealing properties over purely machine-processed data using rules. Second, we describe characteristics of the data: its overdispersion and extreme values or bursts in cyber event counts. We quantify bursts using a Markov model. Bursts are a signature of natural phenomena, including human behavior [21]. Last, we perform temporal prediction of cyber events using the Bayesian State Space Model (BSSM) to predict the number of future events one week ahead. This approach provides both a point estimate and also an interval for the range of forecast uncertainty. Predictive models of analyst-based cyber events may proactively inform CSSPs on a number of important tasks such as resource allocation (e.g., number and location of sensors on the network), estimation of analyst staffing, and the development of effective risk-management strategies.

## Related research

Prior work has demonstrated that cyberattacks, predominantly using machine processed data or alerts, exhibit both deterministic and stochastic patterns. The main finding for deterministic patterns in cyberattacks is that they are neither independent nor random over time. Consequently, the number of attacks in the past helps predict the number of future attacks. The deterministic patterns can often

<sup>1</sup> Note that CSSPs are also be referred to as Computer Defense Service Providers, Computer Network Defense Service Providers, Computer Security Service Providers, Cybersecurity Defense, Managed Computer Security

Service Providers, and Managed Security Service Providers. The last two terms, which included “managed,” explicitly refer to a CSSP that protects the networks and computers for multiple clients/customers.

be leveraged to generate reasonably accurate predictions. When stochastic patterns are present, particularly fluctuations such as extreme values and bursts, temporal forecasting becomes more challenging. The majority of time-series models for count data (i.e., number of attacks) assume a Poisson distribution (equal rate and variance, expressed in a single parameter) [22]. Analysis of automated system data of cyberattacks also reveals systematicity in the tendency of attacks to exploit a disproportionately small set of vulnerabilities.

### Forecasting and bursts

Time-series forecasting has been widely used for prediction of cyberattacks. Using machine-processed attack data from network telescopes and honeypots, the number of cyberattacks over time at minute and hour intervals are predictable over the time period of up to one day [6, 8–10]. Also, cyberattacks were modeled at different levels (attacker IP address, targeted network port[s], and victim IP address) [10]. The use of different levels for attacks were extended to an early warning system by modeling multiple time-series for attack penetration and the number of victims [8]. Other research has reported temporal patterns in attack, but instead used filtered and combined alerts from automated systems [7, 13] or used publically reported cyberattacks from the Hackmageddon database [23]. Last, extreme values—bursts—in attack frequencies have been identified and used to improve model prediction [6, 9–10]. While extreme values in cyberattacks pose a challenge for accurate time-series forecasting, such bursts also underlie human behavior [21].

Bursts in cyberattacks, however, are not a universal pattern. Using analyst verified reports, bursts of cyberattacks were found in only three out of five customer computers/networks protected by a CSSP [12]. Similarly, bursts were not reported for distributed denial-of-service (DDoS) attacks, but these data were limited to one-minute intervals over less than an hour [24]. The mixed finding for the presence of bursts, or lack thereof, raises the question of determining conditions for their occurrence [12]. It is possible that the absence of bursts may be due to insufficient data, the specific method(s) of attack detection, and/or differences in the aggregation/source of the data (e.g., single versus multiple CSSP customers, the type organization for the computers and networks being protected). Nonetheless, when extreme values of attacks are present, and they are not completely random, including them in modeling improves prediction accuracy [6, 9, 10, 12].

### Vulnerabilities

Non-random exploitation of vulnerabilities provides converging evidence for systematicity in cyberattacks. A small number of vulnerabilities tend to comprise the majority of exploits. For example, data breaches can be classified with 90% accuracy using two types of externally observable risks: a) misconfigured internal systems (e.g., not changing the default username and password) and b) anomalous outbound traffic (e.g., spam, port scanning) [25]. Another approach used internal network monitoring logs to identify the probabilities of malware using specific vulnerable vectors (e.g., network configuration, unpatched software, and particular services) [5]. In addition to vulnerability, predictability can be seen in the small number of “attacker” IP addresses, or points of origin, which account for the majority of cyberattacks [4].

### Dataset description

The dataset here comprises 9302 cyber event reports of malware. In general, malware includes computer viruses, worms, trojan horses, malicious mobile code, and blended attacks (combinations of the

previous methods) [26]. Possible attack vectors for delivering malware include email attachments, web browsers (e.g., websites with malicious javascript code or other embedded malicious code), user installed software, and removable media (e.g., flash drives) [26]. The cyber event counts were binned by week, over a time period of  $t = 366$  weeks.

This dataset was obtained from a large DoD CSSP that manages computer and network security for multiple defense agencies and organizations working with the DoD (e.g., industry). The cyber events are occurrences of malware that infected the CSSP customer or clients, not against the CSSP itself. Note that this CSSP rarely relies (<1% of cyber event reports) *only* on automated systems for detecting and documenting potential threats. Human analysts examine the data generated by automated systems and collect sufficient evidence to verify true positives; alert data alone are insufficient without human reasoning or validation. While attack data from automated systems generates orders of magnitude more data than analysts, the “attacks” detected by automated systems include false alarms and typically lack human verification and context. Compared to most previous work on cyberattacks, we likely have higher data quality (fewer false positives) but much lower data quantity. For reasons previously mentioned, the quantity of missed malware is unknown.

In addition, the specific activities performed by CSSPs are highly varied [27–28], so the current dataset may not generalize to all CSSPs. For example, CSSPs can differ in their reliance on automated systems rather than analyst verification. Despite such differences, the majority of CSSPs are responsible for reporting of computer/network security incidents [28]. The generic tasks and workflow for cyber analysts in DoD CSSPs are described in [28] and also [29].

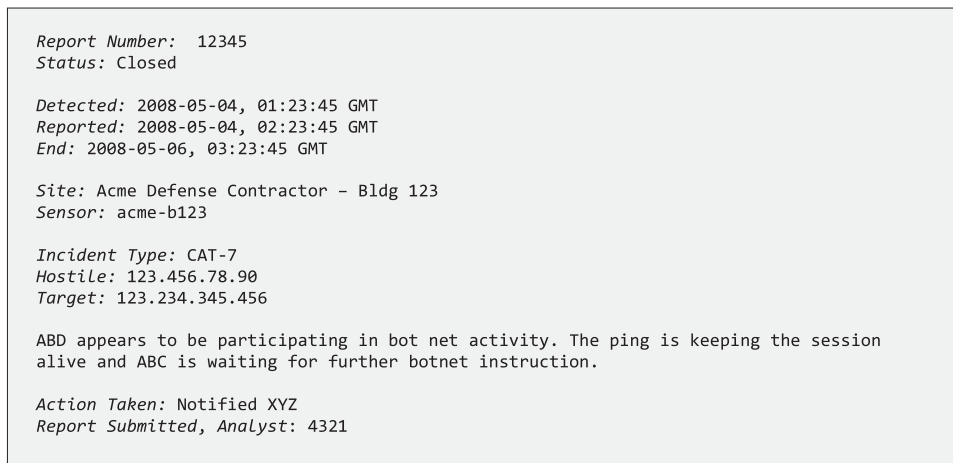
Because the data were from an operational CSSP, we are restricted in providing specific details such as the following:

- i. Actual years for the cyber events (data are from after the year 2000)
- ii. Bins smaller than one week (e.g., events per shift or per day)
- iii. Specific report contents (e.g., method[s] of detection, evidence for the cyber event [also called testimony], type of malware, and infection vector)
- iv. Names of the CSSP customers, the number of customers, and separate analysis by CSSP customer (even if unidentified)

Therefore, the primary analysis examined the frequency of cyber events (counts). The raw cyber event report data was binned into counts of the number of reports generated per week. This interval size is large enough to allow the data and results to be publically released while still providing a large number of data points to show variability and still minimize the number of empty intervals (no events in a week). Use of the data was approved by the US Army Research Laboratory’s Institutional Review Board. Only the aggregated, de-identified dataset and results were approved for public release by operations security for the US Army Research Lab and by the source of the data, the CSSP.

### Cyber events

To provide more background and context for analyst detection we define a cyber event and show a hypothetical cyber event report. A cyber event documents a computer security event and can be shared among other network security operations. This is formally defined as the violation of security, acceptable use, or standard security practices [16]. Although there are specific content requirements for the filing of an incident within the US Federal Government, report content for the CSSP included the date of analyst detection and a



**Figure 1.** Hypothetical example of a cyber event report.

paragraph-long “testimony section” describing evidence for the cyber event. Other information includes the location or site/customer involved and IP addresses for the target and source of the attack. [Figure 1](#) illustrates a hypothetical example of a cyber event report.

An explanation of each of these fields is given to provide greater understanding to the reader.

- Report number—A unique identifier for the incident report
- Status—Indicates if the report is open or closed
- Detected—The initial time that the alert was generated
- Reported—The time that the alert was determined to be significant enough to open a report
- End—Date the issue was resolved
- Site—Location of the attack
- Sensor—Unique identifier of the sensor that captured the alert
- Incident type—Category ranging from 1 to 9, see [30]. Category 7 (CAT-7) (see example above) refers to an incident involving malware or malicious software. Note that we only examine CAT-7 cyber events here.
- Hostile—IP address that issued the command, presumed attacker
- Target—IP address of the host that received the command, defender or potential victim
- Text—Text description of the event(s)
- Action Taken—Describes what action the analyst has taken consistent with CSSP procedures. Typically, the CSSP must notify appropriate POCs at the customer site as well as communicate the findings to peer and superior levels.
- Report Submitted, Analyst—Unique identifier of the analyst who submitted the report

This CSSP has consistent policies and procedures, continuous quality and process monitoring requirements, collaborative reviews, and routine metrics review of how attacks are cataloged. The cyber events that comprise our dataset are all analyst-identified. Each event identified with an analyst report is counted as one instance. Unfortunately, we cannot disclose further details of the counting methodology for security reasons.

## Results

We first summarize characteristics of the dataset: The data were overdispersed and, using a Markov model of burst intensities over

time, we found bursts and bursts of bursts. Next, we describe the forecast model. The best-fit forecast model used a BSSM to predict the number of cyber events for a given week using the number from a previous week (one lag model). Taken together, the main results provide compelling evidence that analyst detected cyber events are not point processes (i.e., random values over time). A secondary result was that the annual increase in cyber events was strongly associated with the rising number of CSSP customers each year.

Results are partially reproducible; the data on counts of weekly cyber events, code to reproduce the figures and models, and the [Supplementary Materials](#) are available here: <https://osf.io/hjffm/>. Unfortunately, for the previously mentioned reasons we are unable to share the raw data on cyber events with content, the code used to clean the raw data, and the data used in the secondary result (annual number of CSSP customers).

## Dataset characteristics

We describe characteristics of the data to summarize it and illustrate the challenges with modeling it. The counts of events were clearly non-normal with overdispersion ([Fig. 2](#)).


To quantify the overdispersion shown in [Fig. 2](#), we use the index of dispersion ( $I$ ), also called the variance to mean ratio, shown in [Equation \(1\)](#) [e.g., 20]:

$$I = \frac{\sigma^2}{\mu} \quad (1)$$

A dispersion index above (1) occurs in data that are overdispersed, because the variance ( $\sigma^2$ ) is greater than the mean ( $\mu$ ) or rate. The index of dispersion in the current data was  $I = 15.29$ , confirming considerable overdispersion. Many time-series models for count data assume a Poisson distribution (equal rate and variance, expressed in a single parameter), and thus dispersion violations may produce incorrect results [see 23].

## Bursts

The presence of bursts is strongly suggested by the overdispersion in the data ([Fig. 2](#) and the index of dispersion). To detect bursts, we leverage the work of Kleinberg [31] using the rates of cyber events (e.g., 1 cyber event per week has a rate of 1, 10 cyber events per week has a rate of 1/10). Using the Kleinberg model, we find bursts in the rates of cyber events starting at the beginning of Year 3 and




**Figure 2.** Shows the histogram of cyber events by week. The x-axis is bins of cyber event counts, and the y-axis depicts the number of weeks in each bin. For example, there were 0–10 cyber events in about 90 of the 366 weeks. The tick marks on the x-axis depict the density of cyber events in each bin and their specific values within the bin; these were randomly jittered by 0.3 to create unique values.

also detect bursts of bursts (Level 2 to Level 3), shown in Fig. 3. Kleinberg bursts use a Markov model that can simultaneously characterize both the normal and anomalous arrival times: multiple states where each state controls the rate of activities. For our dataset, a higher activity state will exhibit lower rates with shorter amounts of time between attacks (more cyber events), whereas a lower activity state will correspond to increased rates because of longer amounts of time between attacks (fewer cyber events). The Kleinberg model switches between the states with a fixed probability that is independent of the state transitions for previous attack rates. For example, bursts are only detected if there are substantial state transitions from higher activity states to lower activity states and vice-versa. The Kleinberg model generalizes this phenomenon by considering all possible rates as an infinite-state Markov model whose parameters can be learnt from the data rather than by using simple thresholds for detecting bursts. Simpler approaches such as models with a limited number of states and thresholding are generally insufficient for burst detection.

Recall our aforementioned point that most cyberattacks exhibit both stochastic and deterministic patterns. If the bursts have a deterministic component, then explicitly modeling them is likely to improve forecast accuracy as others have demonstrated [6, 9–10]. It is important to understand the nature of bursts in the data, as effects of bursts or extreme values may be underestimated in model fits used in forecasting. Because reporting of cyber events was consistent, it is unlikely that the burstiness was an artifact of how events were documented.

### Forecasting

Because of the overdispersion and bursts in the data, the BSSM had better predictive accuracy than the alternative, traditional approaches to time-series forecasting (see [Supplementary Materials](#)). The BSSM had several additional advantages over the other approaches to typical time-series forecasting with regression, such as



**Figure 3.** The x-axis is the year. The y-axis is the burst intensity level, where level 1 is no bursts and higher levels depict more intense bursts. The bars represent the range of burst rates for a given burst intensity level. Note the nesting of bursts with faster and faster rates of time to cyber events for levels 2–3 (increasing burstiness in the number of cyber events). There were a total of 8 weeks with Level 3 bursts and 28 weeks with Level 2 bursts (36 out of 366 weeks exhibited Kleinberg bursts).

an autoregressive (AR) model, autoregressive moving average (ARMA) model, or autoregressive integrated moving average (ARIMA) model. First, the BSSM accommodated overdispersed data using a negative binomial distribution rather than assuming a normal distribution. Second, it can estimate multiple sources of variability (e.g., measurement error) [32] instead of only random model variance [33]. Third, BSSM does not require modeling structural breaks (changes in the forecast at specific time points) and detrending data. Fourth, BSSM can accommodate data with non-stationarity (i.e., changes in mean, variance, and correlation structure over time). State space models have become increasingly popular for forecasting problems because of increased computing power and a family of well-developed Markov Chain Monte Carlo (MCMC) algorithms [34]. At a high level, this approach decomposes the underlying generative process into two types of variables:

- a. Observation variables
- b. State variables.

The model then defines a state transition equation that controls how the process moves between states and an observation equation that generates a noisy output based on the current state.

Forecasting was performed using a BSSM with a negative binomial distribution with a one week lag (the same lag as some of the traditional time-series forecasting methods, described in detail in the [Supplementary data](#)). BSSM is a transparent machine-learning technique that decomposes the data into observations and the model into simultaneously estimated states [32] (see [Equations \(2\) and \(3\)](#) and [Table 1](#)). The BSSM was implemented in the statistical programming language *R* using the “brms package” [35] as a wrapper for the probabilistic programming language *Stan* [36]. Default, non-informative Bayesian priors, also called Empirical Bayes, were used for model estimation. That is, the priors were empirically estimated from the data. [Equation \(2\)](#) is the overall forecast model for observations and [Equation \(3\)](#) is the state; model variables are defined in the [Table 1](#) (the equations and variables are from p. 288 in [37]; also see the specification in [33]).

Forecast (observation):

$$y_t = A_t x_t + v_t, \text{ where } y_t \sim NB(n, p) \tag{2}$$

Local level (state):

$$x_t = \phi x_{t-1} + w \quad (3)$$

No seasonal or cyclical patterns were visually apparent in the data, nor were they meaningful predictors in the alternative models (see the [Supplementary data](#)). Consequently, these parameters were omitted from the BSSM model. Including report length of a cyber event as a parameter in the model yielded a worse fit based on the Widely applicable Bayesian Information Criterion (WBIC) [38]. The model with report length had  $WBIC = 2590.98$ , while the model

omitting report length had  $WBIC = 2586.43$ . Lower WBIC values indicate a better relative model fit.

Figures 4a (weeks for Year 3) and 4b (all 366 weeks) show the one week ahead BSSM forecast. Note that the forecast (white line) in relation to the observed number of events (magenta line). Visible peaks in the magenta line are indicative of bursts in the cyber events. While BSSM leverages deterministic and stochastic patterns, bursts are underestimated in trend lines. Despite deviations in the forecast due to these bursts, nearly all of the observed cyber events were captured by the credible interval (gray-scale shading). However, this assessment of coverage is inference because we do not know the true values of cyber events which is what credible intervals actually estimate. The credible interval (the Bayesian equivalent of a frequentist confidence interval) estimates the probability the true value is captured as a random variable. Formally, the credible interval is defined as the predictive or posterior probability distribution of the model parameters given the observed data [39].

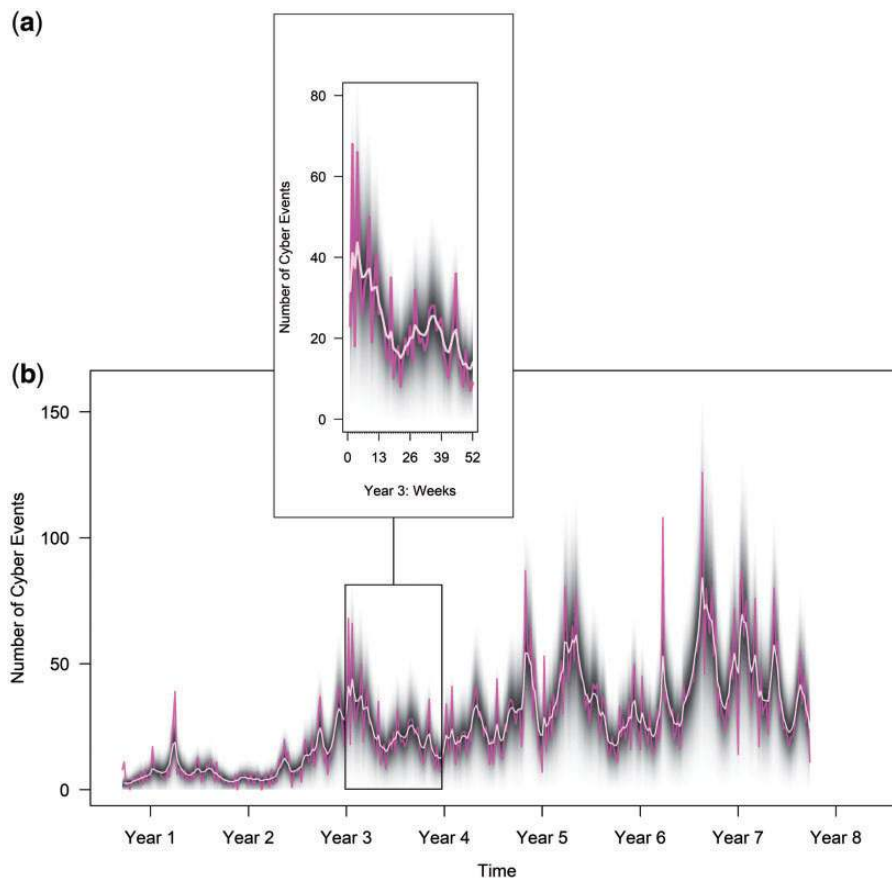
**Table 1.** Definitions of model variables

Variable	Definition
$y$	Estimated number of event(s)
$t$	Time (week)
$NB(n, P)$	$NB$ = negative binomial distribution $n$ = number of failures (zero counts for cyber events) $P$ = probability of success (cyber event occurs)
$\phi$	Latent state
$x_t$	State (transformed known observations from the previous week plus measurement error); note that $x_{t-1}$ is the <i>known</i> observation from the previous week
$A$	Measurement/observation matrix
$v$	Measurement error <sup>a</sup>
$w$	Level error <sup>a</sup>

<sup>a</sup>Note that the credible intervals, described below, incorporates both measurement and level error.

### Forecast accuracy

We quantified the trendline forecast accuracy of the BSSM using multiple measures and compared it to two typical time-series models for all data (Table 2) and for only bursts at Level 3 (Table 3) and Level 2 (Table 4). For all data and every measure of forecast accuracy, the BSSM outperformed the two other models with overall accuracy that subjectively ranged from decent to excellent. The alternative models had similar accuracy to each other for all data as well as both types of bursts. The BSSM also had higher accuracy than alternative models



**Figure 4.** (a) BSSM forecast for weeks in Year 3; see (b) for details. (b) BSSM forecast for all 366 weeks. The forecast trend is the white line, a one-week ahead prediction, and the observed number of events is the magenta line. For the model, given the data, the 95% credible intervals depict uncertainty using gray-scale shading. Gradients from dark to light indicate growing uncertainty in estimating the true number of cyber events, but a higher probability that the interval captures the true value.

for the Level 3 and Level 2 bursts. Nevertheless, the trendline accuracy for the BSSM was subjectively very poor for Level 3 bursts and also poor for Level 2 bursts, both with the exception of single measure of forecast accuracy. As a reminder, Kleinberg bursts were prevalent in about 10% of the data (36 weeks out of 366 weeks). The fairly small number of bursts is a challenge for accurate forecasting, even using a model with a distribution for overdispersed counts. The tradeoffs among different measures of forecast accuracy and equations for each are described by Hyndman and Koehler [40]. Note that the measures of forecast accuracy only reflect the trendline compared to the observed number of cyber events, the true number of cyber events is unknown and can only be estimated.

### Increasing over time: cyber events and CSSP customers

Figure 4 indicates that the overall number of cyber events is increasing over time. However, we found that this closely corresponds to increases in the number of customers per year. Because the specific information is sensitive, we are unable to provide details. We used the number of events per year as the dependent variable and the number of customers per year as the independent variable in ordinary least squares regression. The regression yielded a strong positive slope,  $F(1, 6) = 17.32$ ,  $P < 0.006$ , *Multiple R*<sup>2</sup> = 0.70.

## Summary and discussion

We have shown that analyst-detected and -verified cyber events exhibit sufficient systematicity for time-series forecasting, despite over dispersion. We quantified the presence of bursts and their intensity (~10% data). The low prevalence of bursts, across all CSSP customers, appears consistent with findings of bursts for some CSSP customers [12]. Forecast accuracy for bursts was disappointing, even with the BSSM. Nevertheless, the majority of the data could be forecast with reasonable accuracy. The main finding, that the number of cyberattacks from the previous week helps predict the number that will occur one week later, was strikingly similar to other research using human-verified cyberattacks which was from a university

network [11]. More broadly, the ability to predict cyberattacks over time is consistent with prior research using processed data from network telescopes and honeypots as well as alerts from automated systems [4, 6, 7, 9–13]. In comparison, we used cyber events detected and verified by analysts in an operational CSSP that rarely used automated systems alone and protects critical infrastructure. Like a weather forecast, albeit limited to a one-week ahead prediction here, cyber event forecasting may proactively enhance threat awareness. This may enable CSSPs and similar organizations to better plan for and manage attacks against their defended domains. Moving from reactive and passive defenses to more proactive defenses may help optimize cybersecurity for both analysts and technical systems.

### Potential applications

Advance knowledge about the probable range of attack frequency may aid threat awareness in cybersecurity. Potential applications for CSSPs are using the forecast to proactively inform allocation of capabilities such as the sensors and their configuration (e.g., sampling rate, location of sensors) and type of monitoring (e.g., network traffic). Additionally, a cyber event forecast is an empirical estimation of risk which could be directly applied to models of cyber analyst staffing [41,42].

However, we caution that no forecast should be used as a target or a quota for analysts or CSSPs. When a measure becomes a goal, that measure may no longer be meaningful as an outcome [43]. The meaning of a measure can be distorted by biases such as social and political pressure which may introduce incentives with unintended consequences. Reported detection of “more” attacks may not necessarily improve actual computer and network security.

### Limitations

Because the dataset in this article is unique, analyst detection and verification of cyber events from an operational CSSP with minimal reliance on automated systems, it also comes with several limitations. First, we have only one dataset from one CSSP. This limits the generalizability of the results, although the central findings were

**Table 2.** Measures of forecast trendline accuracy: all data

Measure	BSSM (one week lag)	AR(1) <sup>a</sup>	AR(3)
Mean Absolute Error (MAE)	5.43 cyber events	8.67 cyber events	8.60 cyber events
Mean Absolute Percentage Error (MAPE) <sup>b, c</sup>	68.17% accuracy	48.97% accuracy	49.28% accuracy
Symmetric Mean Absolute Percentage Error (SMAPE) <sup>c</sup>	86.37% accuracy	80.50% accuracy	80.50% accuracy
Root Mean Square Error (RMSE)	8.27 cyber events	13.04 cyber events	12.98 cyber events

AR(1) is a first-order autoregressive model: Time  $t$  is forecast using  $t-1$ , a one week lag. AR(3) is a third-order autoregressive model: Time  $t$  is forecast using  $t-1$ ,  $t-2$ , and  $t-3$ . See [Supplementary Materials](#) for details.

<sup>a</sup>AR(1) is equivalent to ARMA(1, 1).

<sup>b</sup>Accuracy is  $1 - \text{MAPE}$  and  $1 - \text{SMAPE}$ .

<sup>c</sup>MAPE is undefined with actual values of zero because dividing by zero is undefined. Therefore, we excluded the three weeks with zero cyber events and the fitted trend line corresponding to those three weeks.

**Table 3.** Measures of forecast trendline accuracy: level 3 bursts (highest intensity level)

Measure	BSSM (one week lag)	AR(1)	AR(3)
Mean Absolute Error (MAE)	25.46 cyber events	37.81 cyber events	38.05 cyber events
Mean Absolute Percentage Error (MAPE)	42.84% accuracy	34.75% accuracy	33.96% accuracy
Symmetric Mean Absolute Percentage Error (SMAPE)	77.84% accuracy	68.54% accuracy	68.38% accuracy
Root Mean Square Error (RMSE)	30.53 cyber events	47.01 cyber events	47.13 cyber events

**Table 4.** Measures of forecast trendline accuracy: level 2 bursts

Measure	BSSM (one week lag)	AR(1)	AR(3)
Mean Absolute Error (MAE)	12.20 cyber events	13.72 cyber events	13.59 cyber events
Mean Absolute Percentage Error (MAPE)	57.01% accuracy	54.80% accuracy	55.97% accuracy
Symmetric Mean Absolute Percentage Error (SMAPE)	83.41% accuracy	81.71% accuracy	81.75% accuracy
Root Mean Square Error (RMSE)	16.67 cyber events	21.64 cyber events	24.61 cyber events

consistent with prior research. Second, even the best fit model had poor accuracy for predicting bursts. This may be because bursts were relatively rare. Third, we treat each cyber event as equivalent. That is, we do not account for differences in impact among attacks (e.g., affected number of computers and/or networks, consequences on security and economic measures such as loss of productivity and time). Attack severity and attack timing may be related. Recent research indicates that for data breaches, there is a meaningful dependency in the timing between attacks and attack severity (magnitude of the data breach) [44]. Fourth, we could not include the report contents (e.g., method[s] of detection, type of malware) in modeling cyber events and were limited to weekly, rather than finer-grained, counts of cyber events. Such information is sensitive because it could reveal how the CSSP monitors and protects systems. While we were unable to do so here, incorporating internal and external variables is likely to improve the quality of the forecast and also aid in identifying factors relevant to specific attacks.

### Factors for attacks?

Although we can predict malware frequency, we lack direct empirical evidence to explain causes for attack systematicity, as do many others. Moreover, causal inference is a general challenge with observational data [45]. Nonetheless, prediction without identifying causes does not necessarily change the accuracy of prediction. Prediction accuracy will be maintained as long as the conditions and underlying assumptions remain constant. A forecast may become unreliable if conditions change, thus risk models should be frequently recalibrated and validated and should preferably use multiple sources of data [46]. This issue is illustrated by the initial accuracy and then inaccuracy of Google Flu Trends (GFT) [47]. GFT relied on a single source of data (Google search terms related to the flu) and did not update assumptions (e.g., the introduction of suggestions for search terms, other changes to Google search, and media reports) [47].

While we have not identified specific associated or causal factors here, past research suggests that there are multiple causes for cyberattacks. Potential factors are not mutually exclusive. First, there may be planned timing in related attacks (a series of cyberattacks over time) by the same individual or group, or by coordination among groups [48]. In the current work, a series of planned attacks, if they exist, are mixed by aggregation and the absence of detailed information about each event. Recovering separate distributions from their mixtures is challenging [49]. Second, it is possible that exploits are created or purchased on the dark web and deployed by distinct individual or groups around the same time.

Last, activities in social media and events in the physical world likely contribute to attack patterns and vice versa. Prior research has found associations among cyberattacks on DoD networks and foreign media reports of US military actions [50]. Also, website defacements have been linked to a variety of events in the physical world (e.g., violence, protests, and threats) [51]. Incorporating the physical environment as well as expert insights from cybersecurity analysts and analysts in the broader intelligence community may provide

additional predictors for cyberattacks and as well as their associations. Given cyberspace is interdependent with the physical environment, adding predictors from experts and other sources could be used to estimate and model interconnected risks among parameters, see [52–54].

### Future directions

The forecast for analyst detection does not identify specific risk factors associated with attacks. To enhance awareness about specific threats, it is vital to uncover associated and, ideally, causal factors for cyberattacks. This cannot be done with the cyber environment alone because it is inter-dependent with physical environments [1]. In the future, we seek to improve cyber forecasting and to infer the causes for attack patterns. Because of the challenges of openly publishing details with the current dataset, we may use openly available datasets, where attackers sometimes self-identify and even provide the motivation for their attack, such as website defacements (see [51]). Potential variables include events in the physical world as well as more detailed information about the attacks (e.g., the type of malware, exploits/vulnerabilities used, source[s] of the attack, and malware and other cyberattack pricing on the dark web described in [55]). Also, empirically assessing cybersecurity analysts' understanding of the cyber event forecasts could improve its effectiveness for threat awareness. Research on human understanding of uncertainty in visualization of forecast models is surprisingly limited [56].

Another future direction is combining log and automated defenses and prior knowledge of common vulnerabilities with analyst detection. This could advance understanding how layers of defense are coupled, or not, and how particular attacks pass through layers of defense. Network topology is also relevant to attack forecasting: using log and network topology, an early warning system for mitigating attacks has been developed by modeling probable attack penetration and victims [8].

An additional possibility for future research is using the current forecast models and measures of their accuracy as baselines. We are optimistic that others could develop models that have better accuracy than the BSSM forecast accuracy, especially for bursts. Also, our forecast predictions and assessments of their accuracy were limited to one-week ahead. Future work could evaluate the forecast accuracy of h-step ahead (out of sample) using a variety of models.

A final line of future research is to evaluate alternative loss functions. Most time-series forecasting methods assume the squared loss functions for optimization. However, in the security context, especially for critical infrastructure, investigation of other loss functions is a technical gap. For example, a negative forecast error (i.e., underestimate) could be far more expensive than a positive one (i.e., overestimate). The squared loss function treats both scenarios equivalently. Another line of research is to design forecasting models that can produce richer outputs such as range of forecasts along with the confidence interval or other estimation of uncertainty. Additionally, an interactive forecast model that can output the confidence over a human-specified range is often useful from a risk-



aware resource allocation perspective. We also plan to evaluate ensembles of forecasting models so that we can combine the advantages of various forecasting models (e.g., ARIMA, state space-based, and techniques for modeling complex dependencies in the data such as [8]) to produce a superior output.

## Acknowledgments

We thank Bulent Yener, Purush Iyer, and Shouhuai Xu for helpful comments on drafts of the article, and Mark Gatlin for editing the article. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory or the US government. The US government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation.

## Supplementary data

Supplementary data is available at *Journal of Cybersecurity* online.

## Funding

This work was supported by the US Army Research Laboratory Postdoctoral Fellowship Program (E.G.Z.) and Senior Fellowship Program (L.R.M.) administered by the Oak Ridge Associated Universities under Cooperative Agreement Number W911NF-17-2-0003. The work of Gautam Das was supported in part by the Army Research Office/Army Research Laboratory under grant W911NF-15-1-0020 and the National Science Foundation under grant 1745925.

## Conflicts of interest

The authors have no conflicts of interest to declare.

## References

1. Joint Chiefs of Staff. *Cyberspace Operations*. Washington, DC: US Department of Defense, 1–70.
2. Carter A. *The DOD Cyber Strategy*. Washington, DC: US Department of Defense, 2015: 1–42.
3. *DoD Cybersecurity Discipline Implementation Plan*. Department of Defense, 2016.
4. Chen Y-Z, Huang Z-G, Xu S, *et al*. Spatiotemporal patterns and predictability of cyberattacks. *PLoS One* 2015; 10:e0124472.
5. Gil S, Kott A, Barabási A-L. A genetic epidemiology approach to cyber-security. *Sci Rep* 2014; 4:1–7.
6. Peng C, Xu M, Xu S, *et al*. Modeling and predicting extreme cyberattack rates via marked point processes. *J Appl Stat* 2016;44:1–30.
7. Viinikka J, Debar H, Mé L, *et al*. Processing intrusion detection alert aggregates with time series modeling. *Inf Fusion* 2009;10:312–24.
8. Xu M, Hua L, Xu S. A Vine copula model for predicting the effectiveness of cyber defense early-warning. *Technometrics* 2016; 59:508–520.
9. Zhan Z, Xu M, Xu S. Predicting cyberattack rates with extreme values. *IEEE Trans Inf Forensics Secur* 2015; 10:1666–77.
10. Zhan Z, Xu M, Xu S. Characterizing honeypot-captured cyberattacks: statistical framework and case study. *IEEE Trans Inf Forensics Secur* 2013; 8:1775–89.
11. Condon E, He A, Cukier M. Analysis of computer security incident data using time series models. *2008 19th International Symposium on Software Reliability Engineering (ISSRE)* 2008; 77–86.
12. Harang R, Kott A. Burstiness of intrusion detection process: empirical evidence and a modeling approach. *IEEE Trans Inf Forensics Secur* 2017; 12:2348–59.
13. Viinikka J, Debar H, Mé L, *et al*. Time series modeling for IDS alert management. *Proceedings of the 2006 ACM Symposium on Information, Computer and Communications Security*. ACM 2006; 102–13.
14. Irwin B, Riel J-P van. Using InetVis to evaluate Snort and Bro scan detection on a network telescope. *VizSEC Springer, Berlin, Heidelberg*, 2008; 2007:255–73.
15. Patcha A, Park J-M. An overview of anomaly detection techniques: existing solutions and latest technological trends. *Comput Netw* 2007;51: 3448–70.
16. Cichonski P, Millar T, Grance T, *et al*. Computer security incident handling guide. *NIST Spec Publ* 2012;800:61.
17. McHugh J. Testing intrusion detection systems: a critique of the 1998 and 1999 DARPA intrusion detection system evaluations as performed by Lincoln Laboratory. *ACM Trans Inf Syst Secur TISSEC* 2000;3: 262–94.
18. Tavallaee M, Bagheri E, Lu W, *et al*. A detailed analysis of the KDD CUP 99 data set. *IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA)* IEEE, 2009;1–6.
19. KDDCUP 1999 Data. <http://kdd.ics.uci.edu/databases/kddcup99/kddcup99.html> (29 October 2018, date last accessed).
20. Korzyk Sr A. others. A forecasting model for internet security attacks. *National Information System Security Conference*. 1998.
21. Barabasi A-L. The origin of bursts and heavy tails in human dynamics. *Nature* 2005;435:207–11.
22. Jung RC, Tremayne AR. Useful models for time series of counts or simply wrong ones?. *AStA Adv Stat Anal* 2011;95:59–91.
23. Werner G, Yang S, McConky K. *Time Series Forecasting of Cyber Attack Intensity*. Proceedings of the 12th Annual Conference on Cyber and Information Security Research, 2017, 1–3.
24. Fachkha C, Bou-Harb E, Debbabi M. Towards a forecasting model for distributed Denial of Service activities. *IEEE* 2013;110–117.
25. Liu Y, Sarabi A, Zhang J, *et al*. Cloudy with a chance of breach: forecasting cyber security incidents. *24th USENIX Security Symposium (USENIX Security 15)* 2015; 1009–24.
26. Souppaya M, Scarfone K. *Guide to Malware Incident Prevention and Handling for Desktops and Laptops*. National Institute of Standards and Technology, 2013.
27. Killcrece G, Kossakowski K-P, Ruefle RM, *et al*. *State of the practice of computer security incident response teams (CSIRTs)*. 2003.
28. D’Amico A, Whitley K. The real work of computer network defense analysts. In: Goodall JR, Conti G, Ma K-L (eds), *VizSEC 2007*. Berlin Heidelberg: Springer, 2008, 19–37.
29. Cybersecurity Workforce Framework | National Initiative for Cybersecurity Careers and Studies.
30. Chairman of the Joint Chiefs of Staff Manual. *6510.01B: Cyber Incident Handling Program*. Department of Defense 2012.
31. Kleinberg J. Bursty and hierarchical structure in streams. *Data Min Knowl Discov* 2003;7:373–97.
32. Scott SL, Varian HR. Predicting the present with bayesian structural time series. *Int J Math Model Numer Optim* 2014;5:4–23.
33. Hyndman RJ, Koehler AB, Snyder RD, *et al*. A state space framework for automatic forecasting using exponential smoothing methods. *Int J Forecast* 2002;18:439–54.
34. Brooks S, Gelman A, Jones GL, *et al*. (eds.) *Handbook of Markov Chain Monte Carlo*. Boca Raton, FL: Chapman and Hall/CRC, 2011.
35. Bürkner P-C. brms: an R package for Bayesian multilevel models using STAN. *J Stat Softw* 2017;80:1–25.
36. Carpenter B, Gelman A, Hoffman MD, *et al*. Stan: a probabilistic programming language. *J Stat Softw* 2017;76. doi: 10.18637/jss.v076.i01.
37. Shumway RH, Stoffer DS. *Time Series Analysis and Its Applications*. New York, NY: Springer, 2011.
38. Watanabe S. A widely applicable Bayesian information criterion. *J Mach Learn Res* 2013;14:867–97.
39. Kruschke JK, Liddell TM. The Bayesian new statistics: hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective. *Psychon Bull Rev* 2018;25:178–206.
40. Hyndman RJ, Koehler AB. Another look at measures of forecast accuracy. *Int J Forecast* 2006;22:679–88.
41. Ganesan R, Jajodia S, Shah A, *et al*. Dynamic scheduling of cybersecurity analysts for minimizing risk using reinforcement learning. *ACM Trans Intell Syst Technol TIST* 2016;8:1.

42. Ganesan R, Jajodia S, Cam H. Optimal scheduling of cybersecurity analysts for minimizing risk. *ACM Trans Intell Syst Technol TIST* 2017;8:1.
43. Campbell DT. Assessing the impact of planned social change. *Eval Program Plann* 1979;2:67–90.
44. Xu M, Schweitzer KM, Bateman RM, *et al.* Modeling and predicting cyber hacking breaches. *IEEE Trans Inf Forensics Secur* 2018;13: 2856–71.
45. Cox DR, Wermuth N. Causality: a statistical view. *Int Stat Rev* 2007;72: 285–305.
46. Paté-Cornell ME. Uncertainties in risk analysis: six levels of treatment. *Reliab Eng Syst Saf* 1996;54:95–111.
47. Lazer DM, Kennedy R, King G, *et al.* The parable of Google Flu: traps in big data analysis. *Science* 2014;343:1203.
48. Jasper S. *Deterring Malicious Behavior in Cyberspace*. DTIC Document 2015.
49. Lo Y, Mendell NR, Rubin DB. Testing the number of components in a normal mixture. *Biometrika* 2001;88:767–778.
50. Jaros JD. *Determining a Relationship between Foreign News Media Reports Covering U.S. Military Events and Network Incidents Against DoD Networks*. 2005.
51. Sample C. *Cyber + Culture Early Warning Study*. Special report CMU/SEI-2015-SR-025. Retrieved from. [sei.cmu.edu/asset\\_files/SpecialReport/2015\\_003\\_001\\_449739.pdf](http://sei.cmu.edu/asset_files/SpecialReport/2015_003_001_449739.pdf), 2015 (29 October 2018, date last accessed).
52. Szymanski BK, Lin X, Asztalos A, *et al.* Failure dynamics of the global risk network. *Sci Rep* 2015; 5. doi: 10.1038/srep10998.
53. Lin X, Moussawi A, Korniss G, *et al.* Limits of risk predictability in a cascading alternating renewal process model. *Sci Rep* 2017;7. doi: 10.1038/s41598-017-06873-x.
54. Helbing D. Globally networked risks and how to respond. *Nature* 2013; 497:51–59.
55. Nunes E, Diab A, Gunn A, *et al.* Darknet and deepnet mining for proactive cybersecurity threat intelligence. *Intelligence and Security Informatics (ISI), 2016 IEEE Conference On*. IEEE, 2016, 7–12.
56. Ruginski IT, Boone AP, Padilla LM, *et al.* Non-expert interpretations of hurricane forecast uncertainty visualizations. *Spat Cogn Comput* 2016;16:154–72.