

Cyberattack Forecasting using ARIMA with Intensity-based Regressors

Gordon Werner, Ahmet Okutan, Shanchieh Yang, Katie McConky

Rochester Institute of Technology, RIT

Supported by Intelligence Advanced Research Projects Activity (IARPA)



Introduction

- Cyber attacks are an ever increasing threat [1]
- Average cyber breach costs \$3.8 million [2]
- Current defense schemes offer detection of in-progress attacks [3]
- Prevention methods can give victims a warning

[1] Symantec Internet Security Threat Report, <https://www.symantec.com/security-center/threat-report>

[2] IBM Cost of Data Brach Study, <https://www-03.ibm.com/security/data-breach/>

[3] S. Yang, H. Du, J. Holsopple, and M. Sudit. 2014. Attack Projection. In Cyber Defense and Situational Awareness, A. Ko., C. Wang, and R. Erbacher (Eds.). Springer International Publishing, Cham, 239–261. DOI:[h.p://dx.doi.org/10.1007/978-3-319-11391-3_12](https://dx.doi.org/10.1007/978-3-319-11391-3_12)

Predicting Attacks

- Uses previous data trends to predict future behavior
- Daily attack counts have been shown to exhibit correlation
 - Forecast with ARIMA models
- Can models be limited to information taken only from event data?

Motivation

- Construct a forecasting model for cyber incident intensity
- Investigate if a 24 hr. aggregation period is ideal
- Strengthen forecasts using intensity based regressors
- Better understand applicability of ARIMA models to incident prediction
- Provide insightful feedback regarding future intensity to a target
- Explore other forecasting/classification techniques in attack prediction
 - Bayesian networks (BN)


Cyber Incidents as Data

- Analyzing attributes of an incident
- Time of attack
 - Period of the week
- Type of attack
 - Malicious Email, Malicious URL, DOS
- Count of attacks
 - Aggregated over various time periods
- Target of Attack

Cyber Incidents as Data

- Daily cyber incident counts show temporal auto and partial correlation [1]
 - Recent day's volume can indicate future intensity

ACF of Daily E-mail Attack Counts Against Target 2



[1] Werner, Gordon, Shanchieh Yang, and Katie McConky. "Time series forecasting of cyber attack intensity." *Proceedings of the 12th Annual Conference on Cyber and Information Security Research*. ACM, 2017.

Dataset

- 2599 attacks against 2 Targets (*v*):
- Attack Timeline:
 - Target 1. - Jan. 2016 – Oct. 2017
 - Target 2. - Sep. 2016 – Oct. 2017

<i>v</i>	<i>m</i>	Total	Daily Rate	Atk/Day	% of atks
1	Malware	1334	70.2%	2.9	63.3
1	URL	138	15.0%	1.4	6.5
1	E-Mail	636	14.7%	6.5	30.2
2	Malware	169	24.7%	1.7	34.4
2	URL	127	22.0%	1.5	25.9
2	E-Mail	195	33.6%	1.5	39.7

Intensity Forecasting

- Time series of Incident counts:

$X_t = \{x_0, x_1, \dots, x_t\}$, x_i the number of attacks in measurement period, τ

- Predict the number of attacks to occur in next period, x_{t+1}
- Autoregressive moving average (ARMA) model:

$$ARMA(p, q) = \mu + \epsilon + \sum_{i=1}^p \phi_i x_{t-i} - \sum_{i=1}^q \theta_i e_{t-i}$$

- Can generating a time series with different τ improve forecast accuracy?
- Can historical counts aggregated over various τ act as signals for a BN?

Categorical Intensity Forecasts

- Attack count predictions need context
- Can a categorical representation of intensity be forecast with ARIMA?
- Machine Learning approach to classification
 - Bayesian networks

Aggregation and N-Day Ahead

- Can intra-day trends be leveraged to better forecast daily attack count?
 - Predict $\Delta = 24\text{hr}$ with multiple $\tau < 24\text{hr}$ forecasts
- Daily dataset updates are not realistic
 - Predictions need to be made multiple days in advance

Intensity Based Regressors

- Weekly time periods exhibit varying occurrence rates
- Can be used as regressive indicators in ARIMA model

Hour	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
0	0	1	0	1	0	0	0
1	0	1	0	3	0	1	0
2	0	0	0	4	0	0	0
3	0	1	0	1	1	0	0
4	0	0	0	2	1	0	3
5	0	1	2	1	0	1	0
6	0	1	4	1	1	0	1
7	1	0	2	0	4	3	0
8	0	1	1	2	1	1	0
9	1	4	5	3	3	2	0
10	0	5	1	5	3	4	0
11	0	5	5	5	2	1	0
12	0	7	2	0	3	1	0
13	2	3	3	3	2	2	0
14	1	1	0	2	4	1	1
15	1	2	1	3	5	1	0
16	0	2	0	4	0	1	0
17	2	4	2	1	2	0	0
18	0	0	0	3	2	1	0
19	2	0	1	2	1	0	0
20	1	0	1	1	0	0	0
21	0	0	2	0	0	0	1
22	0	0	0	0	0	0	1
23	1	0	2	0	1	0	0
Total:	12	39	34	47	36	20	7

Experimental Baselines

- Intensity prediction baseline
 - Use series mean as forecast
 - Assumes no relationship in the data that can be modeled
- Error Metric
 - Mean absolute error (MAE)

Intensity Prediction Results


- ARIMA able to increase accuracy of predictions in nearly all cases

τ (hrs)	Target 1			Target 2		
	Malware	URL	E-mail	Malware	URL	E-mail
4	07.95%	05.25%	-15.7%	09.30%	09.94%	01.35%
8	05.88%	02.98%	-27.1%	07.30%	08.25%	00.86%
12	04.21%	03.29%	-23.6%	06.81%	11.82%	01.15%
24	05.12%	02.58%	-23.4%	06.74%	14.26%	-0.21%
72	14.80%	00.97%	-33.6%	05.58%	07.45%	-0.59%
120	15.41%	01.15%	-54.2%	02.18%	12.78%	01.91%
168	16.48%	00.19%	-55.2%	05.88%	6.47%	00.00%
336	25.02%	-1.89%	-19.0%	-7.86%	35.10%	03.38%

Table I. ARIMA forecast % improvement over baseline


Intensity Prediction Results

- BN dependency graph relationships correlate to ARIMA results
 - τ leading to better ARIMA prediction show stronger relationships to GT in BN

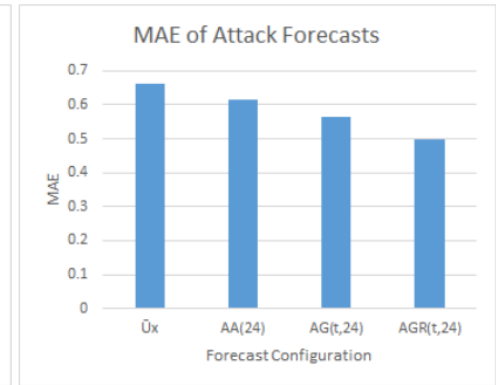


Results (cont.)


- Aggregation outperforms standard ARIMA model for daily predictions
- Regressors not always beneficial
 - No time periods significant enough to improve forecasts




(a) Malware, Target 1




(d) Malware, Target 2



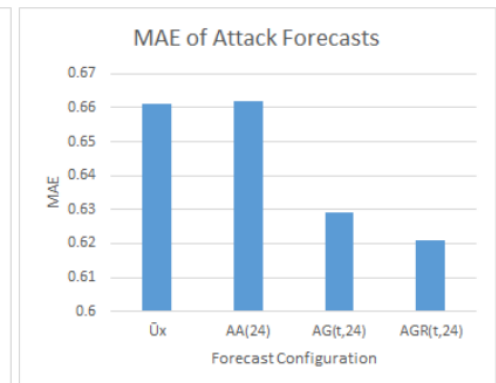
(b) URL, Target 1



(e) URL, Target 2



(c) E-mail, Target 1



(f) E-mail, Target 2

N-Day Ahead Results

- ARIMA models can forecast attack counts accurately up to a week in advance
- Larger N do not have major impact on accuracy

m, v	$AG_N(\tau, 24, 0)$	$AG_N(\tau, 24, 1)$	$AG_N(\tau, 24, 2)$	$AG_N(\tau, 24, 3)$	$AG_N(\tau, 24, 5)$	$AG_N(\tau, 24, 7)$
Malware, T1	18.2%	07.2%	06.7%	06.7%	06.7%	06.7%
URL, T1	10.9%	03.6%	02.6%	01.3%	01.3%	01.3%
E-Mail, T1	-7.3%	-28.8%	-27.7%	-31.2%	-25.2%	-26.3%
Malware, T2	24.9%	06.8%	03.8%	03.9%	04.7%	04.7%
URL, T2	20.6%	15.6%	13.8%	13.8%	12.8%	13.0%
E-Mail, T2	06.1%	00.2%	-00.2%	00.0%	00.5%	00.0%

Table II. N-Day ahead forecast % improvement over baseline

Categorical Forecast Results

- Bayesian networks provide better categorical predictions
- ARIMA sees improvement over baseline

Attack Type	Predictor	AUC
Malware	Naive	.50
	ARIMA	.56
	BN	.61
URL	Naive	.50
	ARIMA	.50
	BN	.50
E-mail	Naive	.53
	ARIMA	.60
	BN	.63

Table III. AUC of categorical predictions

Future Work

- Look at new ways to analyze cyber incidents
 - Arrival process
 - Time between attacks
- Investigate other forecasting methods
 - ARIMA is not necessarily ideal for count series forecasting/classification
 - ARIMA forecasts can be used in conjunction with machine learning techniques
- Expansion of the problem context
 - Investigation of additional regression series
 - External signals
 - Other attack series

Questions?