Program Manager: formerly Dr. Heng Xu now Dr. Sara Kiesler
Theodore Allen allen.515@osu.edu (Principal Investigator)
Gagan Agrawal (Co-Principal Investigator)
Cathy Honghui Xia (Co-Principal Investigator)
Rajiv Ramnath (Former Co-Principal Investigator)

Cyber Vulnerability Maintenance and Optimal Learning

1409214 SBE: TTP Option: Medium: Data-Driven Cyber Vulnerability Maintenance

Key Personnel: Helen Patton, Steve Romig, Solomon Ford
Graduate Students: Chengjun Hou (Ph.D. Candidate, ISE)
Dr. Yue Tan (former, Ph.D. Candidate, ISE)
Dr. Chengjun Hou (former, Ph.D. Student, ISE)
Qiwei Yang (Ph.D Student, CSE), Kaveh Akbari (Ph.D. Student, ISE)
Enhao Liu (M.S. Student, ISE), Tianyu Jiang (M.S. Student, ISE)
Olivia Hernandez

Table 1. NSF 1409214 Data-Driven Cyber Vulnerability Maintenance results.

Challenge	Solution
Base Policy Model – Can we create and implement policies that approximately integrate available scan, incident, and action data?	Use MDP with manual adjustments of transition probabilities, incident probabilities, and average counts (see Jiang, Liu, and Allen, under preparation, and "band aid model" below).
Monitoring Model - How can we monitor to see if there are assignable causes?	Overcome autocorrelation from carried over vulnerabilities, using AR(1) demerit models and effectively chart residuals with simulation-based limits (see below, Afful-Dadzie and Allen, 2016).
Social Media Model – How can we use the power of social media to shed light on vulnerability management?	Focus Tweet streams using Subject Matter Expert Refined Topic (SMERT) models and make manual counts which become observations (Sui, Milam, Allen, 2015 and Allen, Sui, Parker, under review).
Preliminary Software – Can we create a GUI so users can benefit from the base policy model?	Use Visual studio and have file reading and identifications in stage 1 and the knowledge-intensive work is in state 2 (see *Figure 1*).
Model with Optimal Experimentation – Old observations are biased (see below) but how can we plan for and use the new?	Enhance and apply Bayesian Adaptive MDP (BAMDP, Duff 2002) to have observations be simplex points and compound actions to learn many-at-a-time (see Hou, 2015, Allen, Roychowdhury, Hou, under revision).
Model with Enhanced Accuracy from Hosts – Can we use states that are simple and Markovian?	Tree models permit the identification of host features that most accurately predict evolution while permitting implementation (Yang, Allen, Agrawal, 2016)
Model with Enhanced Vulnerability Accuracy and Scan Timing – How can we better model scan, incident, and action data?	Model birth and death of vulnerabilities (state 1). In stage 1, model these. In stage 2, use these models to formulate and solve for host policies (draft early 2017).
Model with Additional Attack Vectors – How can we use near real-time net log data?	Using access both to OSU and ARCYBER net logs, apply discriminant functions and simulation to predict multiple types and costs (draft early 2017).
Real World Applications – Can we demonstrate value in real organization of the associated methods and software?	We have on-going projects with the Ohio State University College of Engineering and Cardinal Health. We have many other customers in mind including Nationwide Insurance and Worthington Cylinders.

Theorem 1. *Consider a BAMDP formulation and a specifically chosen POMDP formulation (details omitted here for conciseness) for any discount factor γ satisfying $0 \le \gamma \le 1$ and any proper observation matrix, $\mathbf{o}^a$ for all $a = 1,\ldots,u$. The BAMMDP formulation and the related POMDP formulations are equivalent such that any feasible solution to one problem is a feasible solution to the other problem. Both solutions have the same objective values and the optimal solution to one problem is the optimal solution to the other problem.*

Table 2. All Windows transition counts (a) actual and (b) counting partial actions as if full actions.

1: Do Nothing	(a) Low	(a) Med. Low	(a) Med. High	(a) High	(a) Critical	(b) Low	(b) Med. Low	(b) Med. High	(b) High	(b) Critical
Low	2,015	52	136	6	0	2,015	52	136	6	(
Med. Low	23	2,379	103	29	31	23	2,379	103	29	3
Med. High	172	77	211,904	1,365	1,910	172	77	211,904	1,365	1,91
High	0	0	0	0	0	**0**	**172**	**77**	**211,904**	**1,36**
Critical	0	0	0	0	0	**0**	**0**	**172**	**77**	**211,90**
2: Research Accept										
Low	0	0	0	0	0	**851**	**115**	**0**	**0**	**(**
Med. Low	0	0	0	0	0	**714**	**851**	**115**	**0**	**(**
Med. High	0	0	0	0	0	**14**	**714**	**851**	**115**	**(**
High	2	14	714	851	115	2	14	714	851	11
Critical	5	5	1,163	129	1,347	5	5	1,163	129	1,34
3: Res. Reject										
Low	0	0	0	0	0	9	0	0	0	(
Med. Low	0	0	0	0	0	6	0	0	0	(
Med. High	0	0	0	0	0	1	384	0	0	(
High	2	14	1,680	0	0	2	14	1,680	0	(
Critical	5	5	1,163	1,476	0	5	5	1,163	1,476	(
4: Compens. Controls										
Low	100	100	0	0	0	100	100	0	0	(
Med. Low	100	100	0	0	0	100	100	0	0	(
Med. High	100	100	0	0	0	100	100	0	0	(
High	100	100	0	0	0	100	100	0	0	(
Critical	100	100	0	0	0	100	100	0	0	(

Table 3. (a) Incident counts by month, (b) compromised host counts by operating system and associated probability estimates, and (c) regression estimates for incident rates.

(a)

Index	Period	Incident Rate
1	1/2014	0.07%
2	2/2014	0.07%
3	4/2014	0.42%
4	5/2014	0.04%
5	6/2014	0.02%
6	7/2014	0.07%
7	8/2014	0.06%
8	9/2014	0.02%
9	10/2014	0.06%
10	11/2014	0.03%
11	12/2014	0.03%
12	1/2015	0.02%
13	2/2015	0.03%
14	3/2015	0.01%
15	4/2015	0.00%
16	5/2015	0.01%
17	6/2015	0.02%
18	7/2015	0.00%
19	8/2015	0.00%
20	9/2015	0.01%
21	10/2015	0.15%

(b)

Rate (Original)	Windows	Linux Enterprise	Other Linux	Other OS
Low	0.49%	2.21%	0.38%	0.05%
Medium Low	0.30%	0.99%	0.74%	0.56%
Medium High	0.03%	1.05%	0.05%	0.05%
High	0.37%	4.90%	0.76%	0.72%
Critical	0.54%	0.00%	0.27%	2.17%

(c)

Rate (Managed)	Windows	Linux Enterprise	Other Linux	Other OS
Low	0.007%	N/A	0.002%	0.004%
Medium Low	0.008%	N/A	0.003%	0.008%
Medium High	0.010%	N/A	0.004%	0.012%
High	0.011%	N/A	0.005%	0.016%
Critical	0.013%	N/A	0.006%	0.020%

Rate (Unmanaged)	Windows	Linux Enterprise	Other Linux	Other OS
Low	0.34%	1.00%	0.425%	0.202%
Medium Low	0.38%	1.80%	0.550%	0.402%
Medium High	0.42%	2.60%	0.675%	0.602%
High	0.46%	3.40%	0.800%	0.802%
Critical	0.50%	4.20%	0.925%	1.002%

Table 4. The base model optimal policy which differs from standard practice (do nothing for low and medium network vulnerabilities and research accept for high and critical network vulnerabilities).

	Windows-normal	Windows-sensitive	Linux Ent.-normal	Linux Ent.-sensitive	Other Linux-normal	Other Linux-sensitive	Other OS-normal	Other OS-sensitive
% of All Hosts	56.68%	2.04%	0.44%	0.05%	10.88%	1.24%	27.71%	0.96%
% Managed %	80%	100%	0%	0%	80%	100%	80%	100%

Unmanaged	Windows-normal	Windows-sensitive	Linux Ent.-normal	Linux Ent.-sensitive	Other Linux-normal	Other Linux-sensitive	Other-normal	Other-sensitive
Low	Do Nothing	Do Nothing	Do Nothing	Do Nothing	Do Nothing	Do Nothing	Do Nothing	Do Nothing
Med. Low	Do Nothing	Do Nothing	Res. Accept	Res. Reject	Do Nothing	Res. Accept	Res. Accept	Res. Accept
Med. High	Do Nothing	Do Nothing	Res. Accept	Res. Reject	Do Nothing	Res. Accept	Res. Accept	Res. Reject
High	Do Nothing	Res. Accept	Res. Accept	Res. Reject	Res. Accept	Res. Accept	Res. Accept	Res. Reject
Critical	Do Nothing	Res. Accept	Res. Accept	Res. Reject	Res. Reject	Res. Reject	Res. Accept	Res. Reject

Managed	Windows-normal	Windows-sensitive	Linux Enterprise-normal	Linux Ent.-sensitive	Other Linux-normal	Other Linux-sensitive	Other-normal	Other-sensitive
Low	Do Nothing	Do Nothing	Res. Accept	Res. Accept	Do Nothing	Do Nothing	Do Nothing	Do Nothing
Med. Low	Do Nothing	Res. Accept	Res. Reject	Res. Reject	Do Nothing	Do Nothing	Res. Accept	Res. Accept
Med. High	Do Nothing	Res. Accept	Res. Reject	Res. Reject	Do Nothing	Do Nothing	Res. Accept	Res. Reject
High	Do Nothing	Res. Accept	Res. Reject	Res. Reject	Do Nothing	Res. Accept	Res. Accept	Res. Reject
Critical	Res. Accept	Res. Accept	Res. Reject	Res. Reject	Do Nothing	Res. Reject	Res. Accept	Res. Reject

Table 5. Proposed TTP work and addressed challenges.

Challenge	Planned Work
Reporting – Can we make useful reports for administrators, managers, and CISOs?	We perform interviews and ethnographic observations to see what decision aids they are using day-to-day and month-to-month. Then, we can enhance these with software and models.
Enhanced Interaction – Can we use developed models to support user decision-making accounting for local and non-local costs and multiple attack models?	Using the enhanced incident data both with simple sums and multicriteria approaches, we illuminate costs born locally from other types of losses and how weightings can support different vulnerability management policies.
Enhanced Experiences – Can we create workflows with improved usability for sustainability at more organizations?	We create a set of workflows that goes beyond the base model, has visualizations including control charts, Twitter Pareto charts, attack model Pareto charts, and time forecasted costs scenarios.
General Applications – Can we improve experiences at many major organizations?	We import and elicit data, generate policies, and assist in their implementation. We measure the costs before and after implementation in net present values.
Elections Applications – Can we improve the policies relating to cyber security in an elections system?	We import and elicit data, generate policies, and assist in their implementation. We measure the costs before and after implementation in net present values. Also, we assist officials to share general best practices in machine allocation, phishing response, and chain of custody.
Machine Learning and Cyber Security Course – Create and enhance a graduate elective on machine learning and cyber security.	We expect to already have proposed a course. If it is approved, we will develop the detailed materials, teach, and assess the course. This will both increase the awareness of cyber issues and MDP and other machine learning technologies.

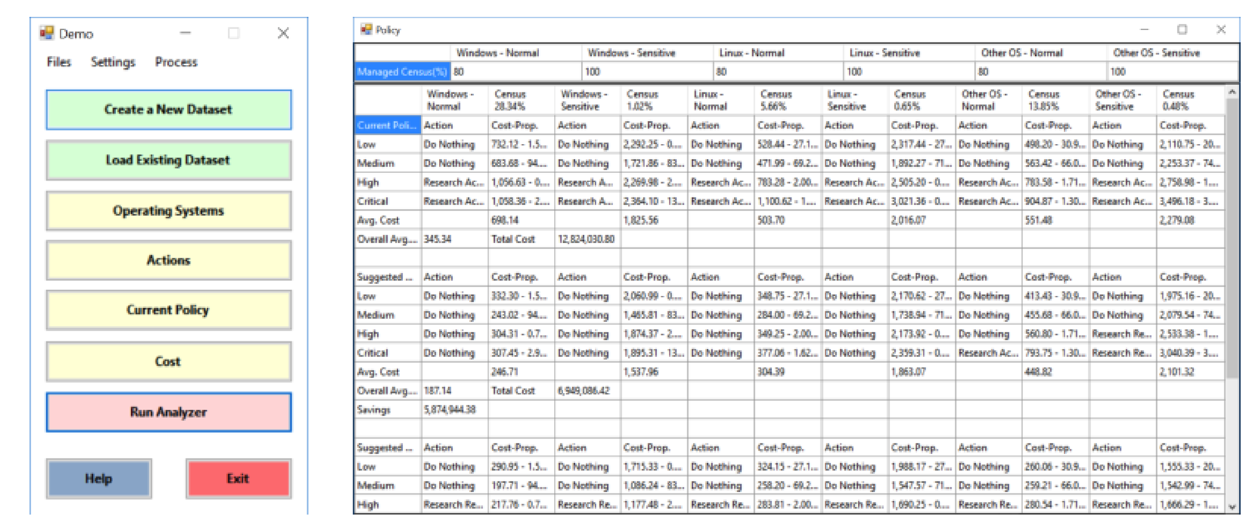

Figure 1. Software related to the base models from NSF project #1409214.

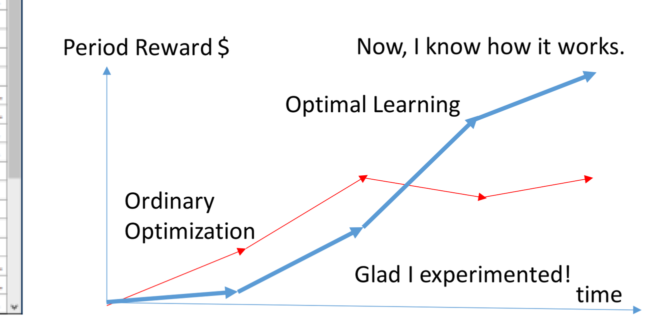

Figure 4. Benefits of experiments with optimal learning.

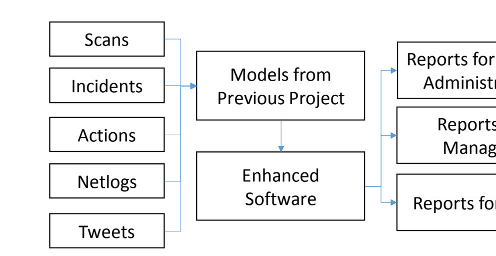

Figure 5. Vision for proposed work.

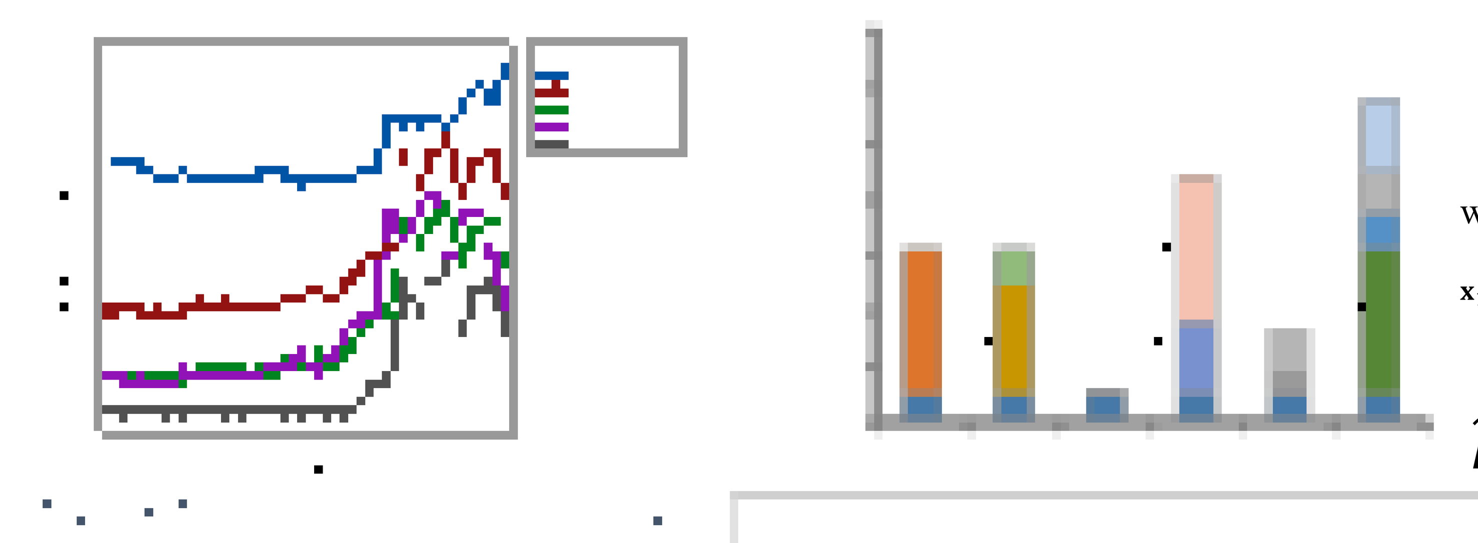

Bayesian Adaptive MDP Formulation

$$Y_t|Y_{t-1}, a_{t-1}, \mathbf{p}_{(k)}^{a_{t-1}}, (k) \sim Multinomial\left[Row_{Y_{t-1}}(\mathbf{p}_{(k)}^{a_{t-1}})\right]$$

where

$$\max_{\mathbf{x}_1,\ldots,\mathbf{x}_{H-1}} \sum_{k=1}^{q} P(k) E_{Y_1,\ldots,Y_H}\left[\sum_{t=1}^{H-1} \gamma^{t-1} r_{Y_t}^{a_t|\mathbf{x}_t} + \gamma^H r_{Y_H}^{a_0|\mathbf{x}_t}\right].$$

$$p_i(y, a_i|O) = \frac{p_0(y,a_i)p(O|y,a_i)}{\sum_{y=1}^{n} p_0(y,a_i)p(O|y,a_i)}$$

Basic MDP Formulation

$C_{action}^{i,OS,a}$ denotes the expected direct cost of taking a specific action a on the hosts starting in state i for a specific OS type in the current period, which is assumed to be period independent.

$C_{incident}^{i,OS}$ denotes the potential cost arising from compromised risk in the hosts being in state i for a specific OS type.

γ is the discount (monthly) factor. Typically, we use 0.99 so that the annual discounting is approximately 10%.

$P_{i,j}^{a}$ refers to the transition probability of a transition from state i to state k under action a.

$Q^{i,OS}$ refers to the probability of an incident for a host in state i and operating system *(OS)*.

$C_{compensation}^{OS,sensitivity}$ is the average cost for incidents on each OS depending on data sensitivity.

The usual value iteration recursion in Markov Decision Processes (MDP) to generate the optimal policy is:

$$\min_{a} C_{total}^{i,OS,a,t} = C_{action}^{i,OS,a} + C_{incident}^{i,OS} + \gamma \sum_{j=1}^{s} P_{i,j}^{a} C_{total}^{j,OS,a,t+1}.$$

Professional Interface For Industrial Applications

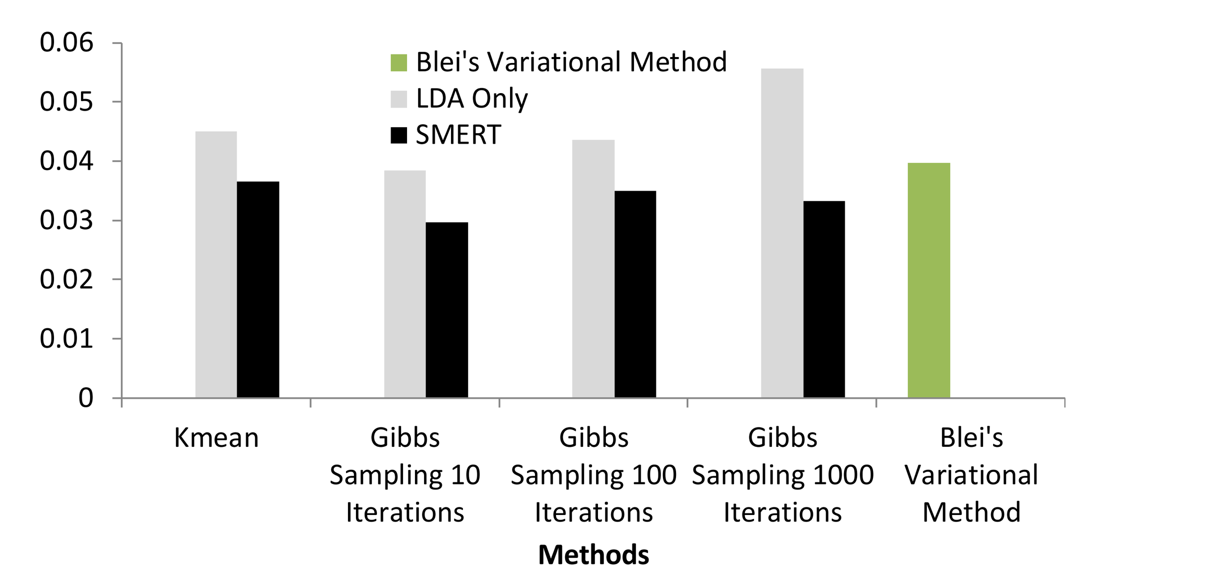

Figure 7. RMS comparison for different estimation methods for Latent Dirichlet Allocation.

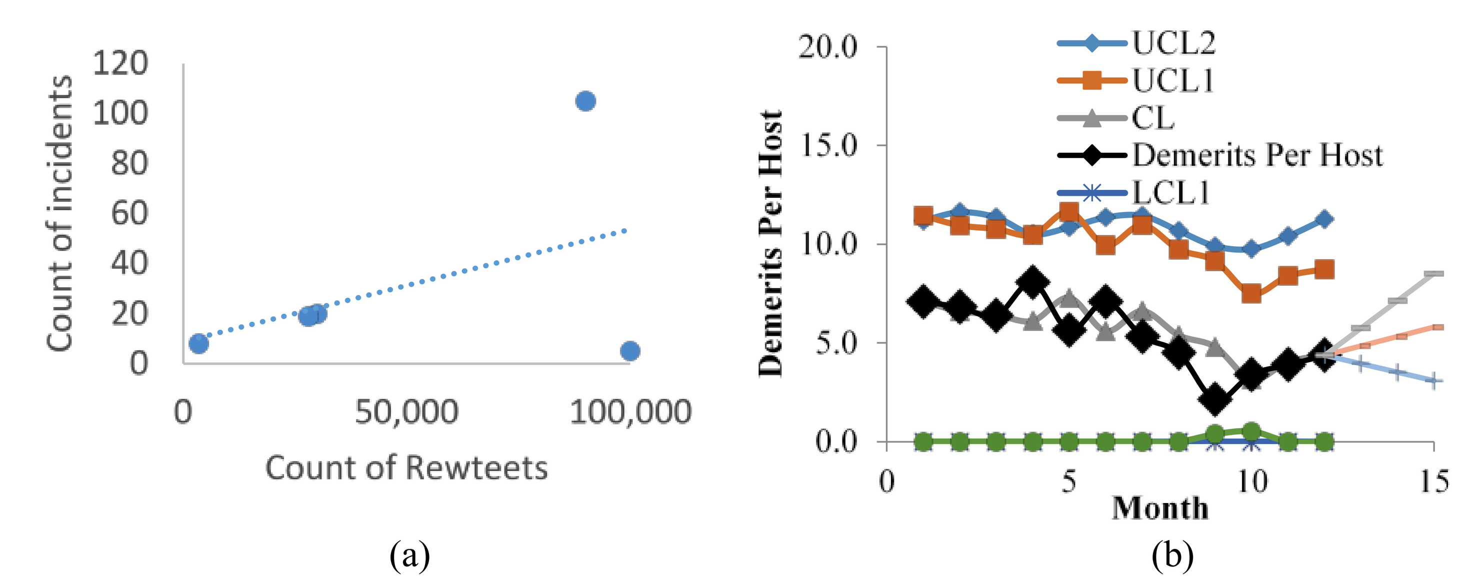

Figure 6. (a) Forecast incidents using retweet counts and (b) forecast total vulnerability demerits.

Plans to Enhance Incident Probability Estimation Using Expanded Rules and Bayesian Adaptive Markov Decision Processes