

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 03-08-2015		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 5-Jun-2011 - 4-Aug-2015	
4. TITLE AND SUBTITLE Final Report: Graph Learning for Anomaly Detection using Psychological Context GLAD-PC			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER W911NF-11-C-0216		
			5c. PROGRAM ELEMENT NUMBER 1M30BM		
6. AUTHORS David Gunning			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Palo Alto Research Center (PARC) 3333 Coyote Hill Road Palo Alto, CA 94304 -1314			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 60135-NS-DRP.17		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT This final report describes the work performed on the GLAD-PC project of the ADAMS program during the final phase of the program (August 1, 2014, July 31, 2015).					
15. SUBJECT TERMS machine learning, data analytics, anomaly detection, insider threat detection, quitting detection					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			David Gunning
					19b. TELEPHONE NUMBER 650-812-4425

Report Title

Final Report: Graph Learning for Anomaly Detection using Psychological Context GLAD-PC

ABSTRACT

This final report describes the work performed on the GLAD-PC project of the ADAMS program during the final phase of the program (August 1, 2014, July 31, 2015).

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received

Paper

08/19/2013 4.00 Jianqiang Shen, Oliver Brdiczka, Yiye Ruan. A comparison study of user behavior on Facebook and Gmail,
ArXiv: 1305.6082, (11 2013): 0. doi: 10.1016/j.chb.2013.06.043

TOTAL: 1

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received

Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Both of these presentations (see attachments) were given at the final ADAMS PI Meeting in Arlington VA, 4-5 MAR 2015:

- (1) Graph Learning and Anomaly Detection using Psychological Context (GLAD-PC)
- (2) Technology transition from PARC

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):**Peer-Reviewed Conference Proceeding publications (other than abstracts):**ReceivedPaper

- 07/14/2015 15.00 Evgeniy Bart, Bob Price, John Hanley. Temporally Coherent Role-Topic Models (TCRTM): deinterlacing overlapping activity patterns,
European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Database. 07-SEP-15, . : ,
- 07/14/2015 14.00 Kumar Sricharan, Gaurang Gavai, Dave Gunning, Rob Rolleston, Mudita Singhal, John Hanley, Juan Julia Liu, Oliver Brdiczka. Detecting employee churn from enterprise social and online activity data ,
2015 ASE Eighth International Conference on Social Computing. 18-AUG-15, . : ,
- 07/14/2015 16.00 Kumar Sricharan, Gaurang Gavai, Dave Gunning, Rob Rolleston, Mudita Singhal, John Hanley, Juan Julia Liu, Oliver Brdiczka3. Detecting insider threat from enterprise social and online activity data,
The 7th ACM CCS International Workshop on Managing Insider Security Threats. 12-OCT-15, . : ,
- 07/21/2014 10.00 Akshay Patil, Juan Liu, Jianqiang Shen, Oliver Brdiczka, Jie Gao, John Hanley. Modeling Attrition in Organizations from Email Communication,
2013 International Conference on Social Computing SOCIALCOM. 08-SEP-13, . : ,
- 08/19/2013 5.00 Jianqiang Shen, Oliver Brdiczka, Juan Liu. Understanding Email Writers: Personality Prediction from Email Messages,
UMAP 2013. 10-JUN-13, . : ,
- 08/19/2013 9.00 Akshay Patil, Juan Liu, Jie Gao. Predicting Group Stability in Online Social Networks,
WWW 2013. 13-MAY-13, . : ,
- 08/19/2013 8.00 Hoda Eldardiry, Evgeniy Bart, Juan Liu, John Hanley, Bob Price, Oliver Brdiczka. Multi-Domain Information Fusion for Insider ThreatDetection,
Workshop on Research for Insider Threat (WRIT). 24-MAY-13, . : ,
- 08/19/2013 7.00 Elise T. Axelrad, Paul J. Sticha , Oliver Brdiczka, Jianqiang Shen. Bayesian network model for predicting insider threats,
Workshop on Research for Insider Threat (WRIT). 24-MAY-13, . : ,
- 08/19/2013 6.00 Francis T. O'Donovan, Connie Fournelle, Steve Gaffigan, Oliver Brdiczka, Jianqiang Shen, Juan Liu, Kendra E. Moore. Characterizing user behavior and information propagation on a social multimedia network,
International IEEE Workshop on Social Multimedia Research (SMMR). 15-JUL-13, . : ,
- 08/20/2012 2.00 Akshay Patil, Juan Liu, Bob Price, Hossam Sharara, Oliver Brdiczka. Modeling Destructive Group Dynamics in On-line Gaming Communities,
International AAAI Conference on Weblogs and Social Media (ICWSM-12). 04-JUN-12, . : ,
- 08/20/2012 3.00 Jianqiang Shen, Oliver Brdiczka, Nicolas Ducheneaut, Nicholas Yee, Bo Begole. Inferring Personality of Online Gamers byFusing Multiple-View Predictions,
Conference on User Modeling, Adaptation and Personalization (UMAP 2012). 16-JUL-12, . : ,
- 08/20/2012 1.00 Oliver Brdiczka, Juan Liu, Bob Price, Jianqiang Shen, Akshay Patil, Richard Chow, Eugene Bart, Nicolas Ducheneaut. Proactive insider threat detection through graph learning and psychological context,
IEEE Workshop on Research for Insider Threat (WRIT). 25-MAY-12, . : ,

TOTAL: 12

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received

Paper

04/01/2015 13.00 J.M. Brown, E.A. Benagh, C.G. Fournelle. Determining Formal and Informal Organizational Hierarchy, International conference on Artificial Intelligence (04 2015)

07/21/2014 12.00 Gaurang Gavai, Sricharan Kumar, Juan Liu, Oliver Brdiczka, John Hanley. Predicting Quitting in the Online Yammer Space, The 8th SNA-KDD workshop (06 2014)

TOTAL: 2

Number of Manuscripts:

Books

Received

Book

TOTAL:

Received

Book Chapter

07/21/2014 11.00 Hoda Eldardiry, Kumar Sricharan, Juan Liu, John Hanley, Robert Price, Oliver Brdiczka, Eugene Bart. Multi-source fusion for anomaly detection:using across-domain and across-time peer-groupconsistency checks, Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA): Innovative Information Science & Technology Research Group (ISYOU), (06 2014)

TOTAL: 1

Patents Submitted

No new patents were submitted in this last year.

Patents Awarded

Awards

None

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	Discipline
Yiran Wang	0.17	
FTE Equivalent:	0.17	
Total Number:	1	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:..... 0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Total Number:

Names of other research staff

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Sri Kumar	0.75
Gaurang Gavi	0.75
David Gunning	0.20
John Hanley	0.25
Mudita Singhal	0.10
Eugene Bart	0.10
Bob Price	0.10
FTE Equivalent:	2.25
Total Number:	7

Sub Contractors (DD882)

1 a. Boston Fusion

1 b. 1 Van de Graaff Dr.

Ste 107

Burlington

MA

01803

Sub Contractor Numbers (c):

Patent Clause Number (d-1):

Patent Date (d-2):

Work Description (e): Boston Fusion will provide data analysis support.

Sub Contract Award Date (f-1): 7/1/13 12:00AM

Sub Contract Est Completion Date(f-2): 6/4/15 12:00AM

Inventions (DD882)

Scientific Progress

(1) Foreword

This report details the scientific progress of PARC team in the GLAD-PC project for the period of 08/01/2013- 07/31/2014. This research has been funded by the DARPA/ADAMS program under contract W911NF-11-C-0216. Any opinions, findings, and conclusions or recommendations in this report and associated material are those of the authors and do not necessarily reflect the views of the governmentfunding agency.

(2) Table of Contents (None)

(3) List of Appendixes:

- Detecting insider threat from enterprise social and online activity data
- Temporally Coherent Role-Topic Models (TCRTM): deinterlacing overlapping activity patterns
- Detecting employee churn from enterprise social and online activity data
- PARC-ADAMS-PI-Meeting-20150305_v4

(4) Statement of the problems studied

The PARC team investigated three approaches to detecting aspects of malicious insider activity: a) psychological profiling from email; b) quitting dynamicsand quitting prediction from corporate social media data; and c) detecting unusual and anomalous behavior from on-line activities.

(5) Summary of the most important results

With regard to (a) Psychological profiling from email: we have defined a Bayesian model for the motivations and psychology of the malicious insider and an associated degree of interest. We aimed then to predict the derived psychological variables automatically from text in emails. Several large studies have been conducted involving over 1000 subjects. We measured the subjects' psychology using surveys and collected anonymized features from their email communications. We were able to predict the subjects' psychological variables with up to 95% accuracy (see [Shen1]). The constructed predictors have been applied to various real-world data sets including large corporate email data sets. The results have been made accessible to analysts via a specific personality prediction visualization called the Interactive Personality Workbench (described in last years AUG 2013 – JUL 2014 Interim Report). Initial feedback we received from the analysts is very positive.

With regard to (b) quitting dynamics and quitting prediction from corporate social media data. Last year, we have looked into predicting if and when people quit a corporation using their activity on an internal social media network called Yammer. We got access to a data set of over 24,000 corporate users of this internal social media network of a large corporation, including over 2,000 groups and over 150,000 public messages. The goal was to predict, at any given time instance, if an employee is likely to quit the company. For quitting the company, we have identified 298 quitter instances among 7000 non-quitter instances (after cleaning and filtering the data set according to appropriate parameters, e.g. number of messages and activity scores). Using a random forest and a balanced data sets (50% baseline), we get an accuracy of 68%, which means an improvement of 36% over the baseline. A detailed summary of the results including figures and tables can be in [Gava1].

During this last year we extended this work quitting dynamics by studying employee churn behavior. Employee churn is a significant concern for organizations, with downsides including loss of talent, its productivity, and also security risk, given that employees are likely to retain confidential company data after they quit. PARC developed hypothesizes that precursors to an employee quitting a company will manifest in the enterprise social and online activity data of the employee. To this end, we processed and extracted relevant features from social data including email communication patterns and content, and online activity data such as web browsing patterns, email frequency, and file and machine access patterns, and used these features to build a predictive model for detecting employee quitting events ahead of time. We tested our predictive models on two different real world data sets, and our experiments show that we are able to detect quitting events with moderately high accuracy. Finally, we build a visualization dashboard that enables managers and HR personnel to quickly identify employees with high quitting scores, which will enable them to take suitable preventive measures to reduce, churn [Sricharan2, attached].

Regarding (c) detecting unusual and anomalous behavior from on-line activities, PARC investigated techniques to discover insider threat in organizations by identifying abnormal behavior in enterprise social and online activity data of employees. To this end, we processed and extracted relevant features that were possibly indicative of insider threat behavior. This includes features extracted from social data including email communication patterns and content, and online activity data such as web browsing patterns, email frequency, and file and machine access patterns. Subsequently, we detect statistically abnormal behavior with respect to these features using state-of-the-art anomaly detection methods, and declare this abnormal behavior as a proxy for insider threat activity. We tested our approach on a real world data set (the Vegas data set from ADAMS) with artificially injected insider threat events. Our experiments show that our proposed approach is fairly successful in identifying insider threat events. Finally, we build a visualization dashboard that enables managers and HR personnel to quickly identify employees with high threat risk scores, which will enable them to take suitable preventive measures and limit security risk [Sricharan1, attached].

PARC also investigated the specific problem of identifying overlapping activity patterns in the VEGAS data set. The Temporally Coherent Role-Topic Model (TCRTM) is a probabilistic graphical model for analyzing overlapping, loosely temporally structured activities in heterogeneous populations. Such loose temporal structure appears in many domains, but especially in the ADAMS data, where individual events that make up an activity have coherence, but no strong temporal ordering. For instance, preparing a PowerPoint presentation may involve opening files, typing text, downloading images, and saving files. These activities occur together in time, but without a strong ordering or fixed duration. These temporally coherent activities may also overlap – the user might also be responding to email while working on the presentation. Finally, the population of users has subgroups – in the office, administrators, salespeople and engineers will have different activity distributions. The unique architecture of the TCRTM model allows it to automatically infer an appropriate set of roles and activity types while simultaneously assigning users to these roles and segmenting their event streams into high-level activity instance descriptions. On two real-world datasets taken from computer user monitoring and social services debit card transactions we show that TCRTM extracts semantically meaningful structure and improves perplexity score on hold-out data by a factor of five compared to standard models such as LDA [Bart1, attached].

All of these results and summary of PARCs work on ADAMS was presented at the final ADAMS PI Meeting, held at DARPA, in March 2015 (the briefing slides are attached).

(6) Bibliography

Books

Hoda Eldardiry, Kumar Sricharan, Juan Liu, John Hanley, Robert Price, Oliver Brdiczka, Eugene Bart. Multi-source fusion for anomaly detection: using across-domain and across-time peer-group consistency checks, *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA): Innovative Information Science & Technology Research Group (ISYOU)*, (06 2014)

Peer-Reviewed Conference Proceeding publications (other than abstracts)

Kumar Sricharan, Gaurang Gavai, Dave Gunning, Rob Rolleston, Mudita Singhal, and John Hanley, Detecting insider threat from enterprise social and online activity data, the 7th ACM CCS International Workshop on Managing Insider Security Threats, 12-16 OCT 2015.

Evgeniy Bart, Bob Price, and John Hanley, Temporally Coherent Role-Topic Models (TCRTM): deinterlacing overlapping activity patterns, *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 7-11 SEP 2015.

Kumar Sricharan, Gaurang Gavai, Dave Gunning, Rob Rolleston, Mudita Singhal, John Hanley, Juan Julia Liu, and Oliver Brdiczka, Detecting employee churn from enterprise social and online activity data, 2015 ASE Eighth International Conference on Social Computing, 18-20 AUG 2015.

Akshay Patil, Juan Liu, Jianqiang Shen, Oliver Brdiczka, Jie Gao, John Hanley. Modeling Attrition in Organizations from Email Communication, 2013 International Conference on Social Computing SOCIALCOM. 08-SEP-13, . . .

Akshay Patil, Juan Liu, Jie Gao. Predicting Group Stability in Online Social Networks, WWW 2013. 13-MAY-13, . . .

Hoda Eldardiry, Evgeniy Bart, Juan Liu, John Hanley, Bob Price, Oliver Brdiczka. Multi-Domain Information Fusion for Insider Threat Detection, Workshop on Research for Insider Threat (WRIT). 24-MAY-13, . . .

Elise T. Axelrad, Paul J. Sticha, Oliver Brdiczka, Jianqiang Shen. Bayesian network model for predicting insider threats, Workshop on Research for Insider Threat (WRIT). 24-MAY-13, . . .

Francis T. O'Donovan, Connie Fournelle, Steve Gaffigan, Oliver Brdiczka, Jianqiang Shen, Juan Liu, Kendra E. Moore. Characterizing user behavior and information propagation on a social multimedia network, *International IEEE Workshop on Social Multimedia Research (SMMR)*. 15-JUL-13, . . .

Jianqiang Shen, Oliver Brdiczka, Juan Liu. Understanding Email Writers: Personality Prediction from Email Messages, UMAP 2013. 10-JUN-13, . . .

Jianqiang Shen, Oliver Brdiczka, Nicolas Ducheneaut, Nicholas Yee, Bo Begole. Inferring Personality of Online Gamers by Fusing Multiple-View Predictions, *Conference on User Modeling, Adaptation and Personalization (UMAP 2012)*. 16-JUL-12, . . .

Akshay Patil, Juan Liu, Bob Price, Hossam Sharara, Oliver Brdiczka. Modeling Destructive Group Dynamics in On-line Gaming Communities, *International AAAI Conference on Weblogs and Social Media (ICWSM-12)*. 04-JUN-12, . . .

Oliver Brdiczka, Juan Liu, Bob Price, Jianqiang Shen, Akshay Patil, Richard Chow, Eugene Bart, Nicolas Ducheneaut. Proactive insider threat detection through graph learning and psychological context, *IEEE Workshop on Research for Insider Threat (WRIT)*. 25-MAY-12, . . .

Manuscripts submitted, but not published

Gaurang Gavai, Sricharan Kumar, Juan Liu, Oliver Brdiczka, John Hanley. Predicting Quitting in the Online Yammer Space, *The 8th SNA-KDD workshop* (06 2014)

Papers published in peer-reviewed journals


Jianqiang Shen, Oliver Brdiczka, Yiye Ruan. A comparison study of user behavior on Facebook and Gmail, *ArXiv: 1305.6082*, (11 2013): 0. doi: 10.1016/j.chb.2013.06.043

Technology Transfer

See attachment: Technology transition from PARC

“Graph Learning and Anomaly Detection using Psychological Context (GLAD-PC)

The PARC team, 03/05/2015



The Team



David Gunning
PI



John Hanley
Info. Architecture
Database



Mudita Singhal
Visualization
Design



Oliver Brdiczka
Quitting Example



Sricharan Kumar
Machine learning



Gaurang Gavai
Info. Architecture
Feature Extraction



Rob Rolleston
Visualization Design
& Implementation



Julia Liu
Quitting Example

Subcontractors



Boston Fusion
Remote operations, group dynamics analysis

Betsey Benagh

Joanna Brown

Connie Fournelle

Kendra Moore


Personality and Anomalous Behavior



- Organizations (and society) face increasing amount of threats from “inside” and “outside”
- Challenge: Uncover malicious behavior in a *timely* way through automatic analysis
- Anomalous behavior trace often precedes the actual “incident”
- Personality has been shown to be a reliable indicator for future (malicious) behavior (Jaclyn et al., 2011)
- 50% of job quitters steal confidential company data

ADAMS GLAD-PC

- PARC project:
 - Graph Learning and Anomaly Detection using Psychological Context (GLAD-PC)
 - **Idea:** combine *graph learning / structural anomaly detection* and *psychological modeling*



Previous Research: Personality Profiling for Malicious Insider Detection

- We are interested in **psychological profiles** as a indicators for future malicious behavior
- Why ?
 - Counterproductive (cyber-)behaviors have been shown to be **highly correlated** with Big-5 personality variables ^[1]
 - Actual insider threats have a low base rate → psychological profiles are a powerful filter to reduce false positives

[1] Jaclyn M. Jensen, Pankaj C. Patel, Predicting counterproductive work behavior from the interaction of personality traits, Personality and Individual Differences 51(4):466-471, Sept. 2011.

What is a Personality Profile?

Personality Variables

Neuroticism

Agreeableness

Conscientiousness

Excitement Seeking

Hostility

Extraversion


Self-Assurance

Overall Mood/Emotion

Organizational Deviance

Personal Deviance


Perceived Stress



Degree of Interest



Approach

- Idea: automatically estimate personality from emails



Results

- Data Collection & Evaluation Results (of Estimators):
 - Over 1000 personality profiles + emails collected from MechTurk and company-internal (for training the estimations)




Interactive Personality Visualization

User: user0010 ▾

x-Axis: Hostility ▾ y-Axis: Perceived Stress ▾

Agreeableness	Conscientiousness	Excitement Seeking	Extraversion	Hostility	Interpersonal Deviance	Neuroticism	Overall Mood	Perceived Stress	Self Assurance
13.5	30	30.5	32.8	26.5	49.8	46.1	13.9	28.2	49.8



Previous Research: Quitting and Destructive Group Dynamics

- We are interested in **quitting behavior & destructive group dynamics**
- Proxy of malicious behavior: “50% job leavers steal confidential company data”
- Questions:
 - Can we observe quitting behavior and destructive group dynamics in **real-world** and **social space**
 - How is real-world behavior related to social space data
 - Can we predict real-world behavior

Previous Research on Quitting Behavior

Online Games



Destructive group dynamics


- if/when a player will quit a guild
- damage associated with a quit event
- guild stability against member loss

Startup Venture


Churn prediction in a real-world corporation



Predict quitting based on work practice, email, and content.




Yammer



2014 Summer Intern - Yiran Wang

Literature Search



Structured Interviews

- Recent quitters
- N=12 (Male = 9, Female = 3)
- Job titles include:
 - research scientist (2)
 - software engineer (3)
 - research engineer (2)
 - director/manager (1)
 - senior associate in a bank (1)
 - system engineer (1)
 - manufactory engineer (1)
 - office manager (1)


Interview Results

- Web Browsing
 - Increased use of career sties (e.g., LinkedIn)
 - Increased browsing of company profiles
- Personal Email
 - Increased use of personal email for job applications
- Work Email
 - No conscious change
 - Some made an effort to maintain normal email behavior
- Work Routine
 - Shortened work hours and more time off to accommodate interviews
- Multitasking
 - Shortened attention spans at work
 - more task switching
- Engagement
 - Decreasing engagement in general
 - More neutral sentiment in emails

Features Selection/Engineering

- Extract a rich set of features:
 - Email Usage (-sent count)
 - Email Content (-subject char length)
 - Log On / Log Off Statistics
 - Application Activity (+max time spent on activity, + # of activity types per day)
 - Web Usage (time on –internal/+job sites)
 - Feature matrix F : $U \times T \times D$

This Year's Problem Set-up



Problem set-up

- Twin approaches:
 - Supervised – Use quitting labels as proxy
 - Build classifier to predict quitters and corresponding time instances
 - Unsupervised – Use anomaly detection methods to detect abnormal behavior

Vegas Dataset

- Multi-Domain Employee Data
- Anonymized application-wise log of User activity
- Anonymized activity log of user interactions with different agents
- Email interaction data between business unit users
- Aggregated statistics on Email content data
- Snapshots of LDAP hierarchy
- ~~Day-to-day LDAP diffs~~

Vegas Dataset

Dataset Statistics	
Date range	2013-10-01 to 2014-07-01 (8 months)
Users	6805 users
Dataset Size	~ 1 billion User Activity Records
Domains	Email Usage, Email Content, Logon Logoff, Application Usage, Web Usage
Target Users	- 555 Quitters (1270 Pseudo) - 104 Red Team Users

Feature Extraction

- Calculated aggregate features from raw data
- Constructed features in 5 different domains
- Features developed from earlier Yammer work were supplemented with newer ones derived from insights gained by conducting interviews with employees that quit their jobs


Email Usage Features	Logon Logoff Features
Weekly sent count	Number of logons
Weekly read count	Number of logoffs
Number of messages sent in the day	Number of hours with logon activity
Number of messages sent at night	Number of hours with logoff activity
Number of messages read in the day	Activity Features
Number of messages read at night	Number of activity types
Email Content Features	Max contiguous time spent on activity
Average subject word length	Number of activities
Average subject character length	Time spent on on email applications
Average content character length	Time spent on on productivity applications
Average content word length	Time spent on on web applications
Average content sent length	Time spent on on engineering applications
Number of exclamation points	Web Usage Features
Number of multiple exclamation points	Time spent on on websites
Number of question marks	Time spent on on career sites
Number of multiple question marks	Time spent on on web mail sites
Number of brackets	Time spent on on entertainment sites
Number of dashes	Time spent on on internal SM sites
Number of double dashes	Time spent on on internal sites
Number of ellipses	Time spent on on news sites
Number of commas	Time spent on on private social media sites
Number of semicolons	Time spent on on search sites
Number of colons	Time spent on on tech sites

Hierarchy Creation

- Needed a hierarchy of the organization to be able to compare the behavior of a user with their peers
- Data available: daily snapshots of LDAP hierarchy
- We created a normalized hierarchy by finding the most persistent relationships between supervisors and employees over the time period in consideration
- Resulted in ~200 sub-trees due to the business unit not containing the higher levels of the hierarchy

Hierarchies

- Examples of sub-trees
- ext---- nodes are external to the business unit



Supervised approach

Quitting Detection

Problem statement: At any given time, predict if an employee is likely to quit the company:

- Restrict attention to (User U, Time T) tuples such that user U has data for at least 1 month leading up to time T
- 0.6M such total instances; 2K / 0.6M ($\sim 0.5\%$) instances are when user U has quit in time T, T-1 or T-2
- Subsample to deal with class-imbalance problem

Supervised approach

Quitting Detection


- Accuracy = 73% using Random forests
(46% improvement compared to random baseline)
- Content features are most predictive for quitters and pseudo-quitters

- Confusion Matrix:

Class	+	-
+	0.746	0.254
-	0.310	0.690

Supervised approach


Quitting Detection




Quitting Visualization Dashboard

ADAMS Dashboard

Quitting Risk and Anomaly Detection Features Using 4,524 out of 4,524 records | [Reset All](#)



Reference:




Average values over date range: **2013-10-07** to **2014-06-30**

[Show / hide columns](#)

user6	Quitter	Pseudo	RedTeam	Score
06F2A3	true	true	false	0.21
142F9E	true	true	false	0.10

user6	Quitter	Pseudo	RedTeam	Score
-------	---------	--------	---------	-------

Quitting Visualization Dashboard



Supervised approach

Quitting Detection - Insight


- Quitting scores tend to peak ~2 weeks before quitting



Supervised approach

Quitting Detection - Insight

- Quitting scores tend to peak ~2 weeks before quitting



Unsupervised approach

Quitting Detection

Detect anomalies with respect to two aspects:

- Detect if user is anomalous with respect to rest of the employees at each time instance
- Detect if user's behavior has changed drastically over time
- Idea: In addition to features F , also construct differences


$$dF = F[:, T+1, :] - F[:, T, :]$$

- Run iForest on joint matrix $[F; dF]$


Unsupervised approach

Quitting Detection

Ranks of Red Team Users




Ranks of Non Red Team Users



Unsupervised approach

Quitting Detection

Ranks of Red Team Users






- Can identify 46% of red-team events by tracking top 15% of users every week
- 85% by tracking top 35%

Anomaly Visualization Dashboard

ADAMS Dashboard

Quitting Risk and Anomaly Detection Features Using 4,524 out of 4,524 records | [Reset All](#)




Average values over date range: 2013-10-07 to 2014-06-30

[Show / hide columns](#)

user6	Quitter	Pseudo	RedTeam	Score
03C576	false	false	true	0.44
3671A4	false	false	true	0.46
7E75F8	false	false	true	0.12

Anomaly Visualization Dashboard




Conclusion and Future Work

- End-user activity can be used to determine suspect insider threat behavior
- False-alarms fairly significant, due to
 - Rarity of abnormal events
 - Statistical anomalies that do not translate to real world

Conclusion and Future Work

- Further research needed to bring down false alarm rate
 - Integration of external data sources
 - Integration of psychological modeling
 - Incorporating analyst feedback to select features



Thank you!

Questions?

