

Beyond The Hype:

Machine Learning & AI for Security Operations

Anthony Tellez, CISSP, CEH, CNDA

Staff Data Scientist- Machine Learning & AI | Splunk

June 2019



Forward-Looking Statements

During the course of this presentation, we may make forward-looking statements regarding future events or the expected performance of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results could differ materially. For important factors that may cause actual results to differ from those contained in our forward-looking statements, please review our filings with the SEC.

The forward-looking statements made in this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, this presentation may not contain current or accurate information. We do not assume any obligation to update any forward-looking statements we may make. In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only and shall not be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionality described or to include any such feature or functionality in a future release.

Splunk, Splunk>, Listen to Your Data, The Engine for Machine Data, Splunk Cloud, Splunk Light and SPL are trademarks and registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names, or trademarks belong to their respective owners. © 2017 Splunk Inc. All rights reserved.

Intro - Anthony Tellez

CISSP, CEH, CNDA, Sec+


- ▶ | where _time@Splunk > 5y
- ▶ Previous:
 - U.S. Gov Contractor, Geospatial Analyst
- ▶ Specializations
 - Cryptography
 - Information Security – Red Team
- ▶ Data Scientist
 - Security & Fraud Analytics
 - Data Visualization & Statistics
- ▶ Responsible for the relationship between emerging technologies and field organization
 - Acquisitions
 - Incubation
 - Product Development
- ▶ <https://github.com/anthonygtellez/>
- ▶ Fact: Spends 80% of the year on a plane traveling globally.

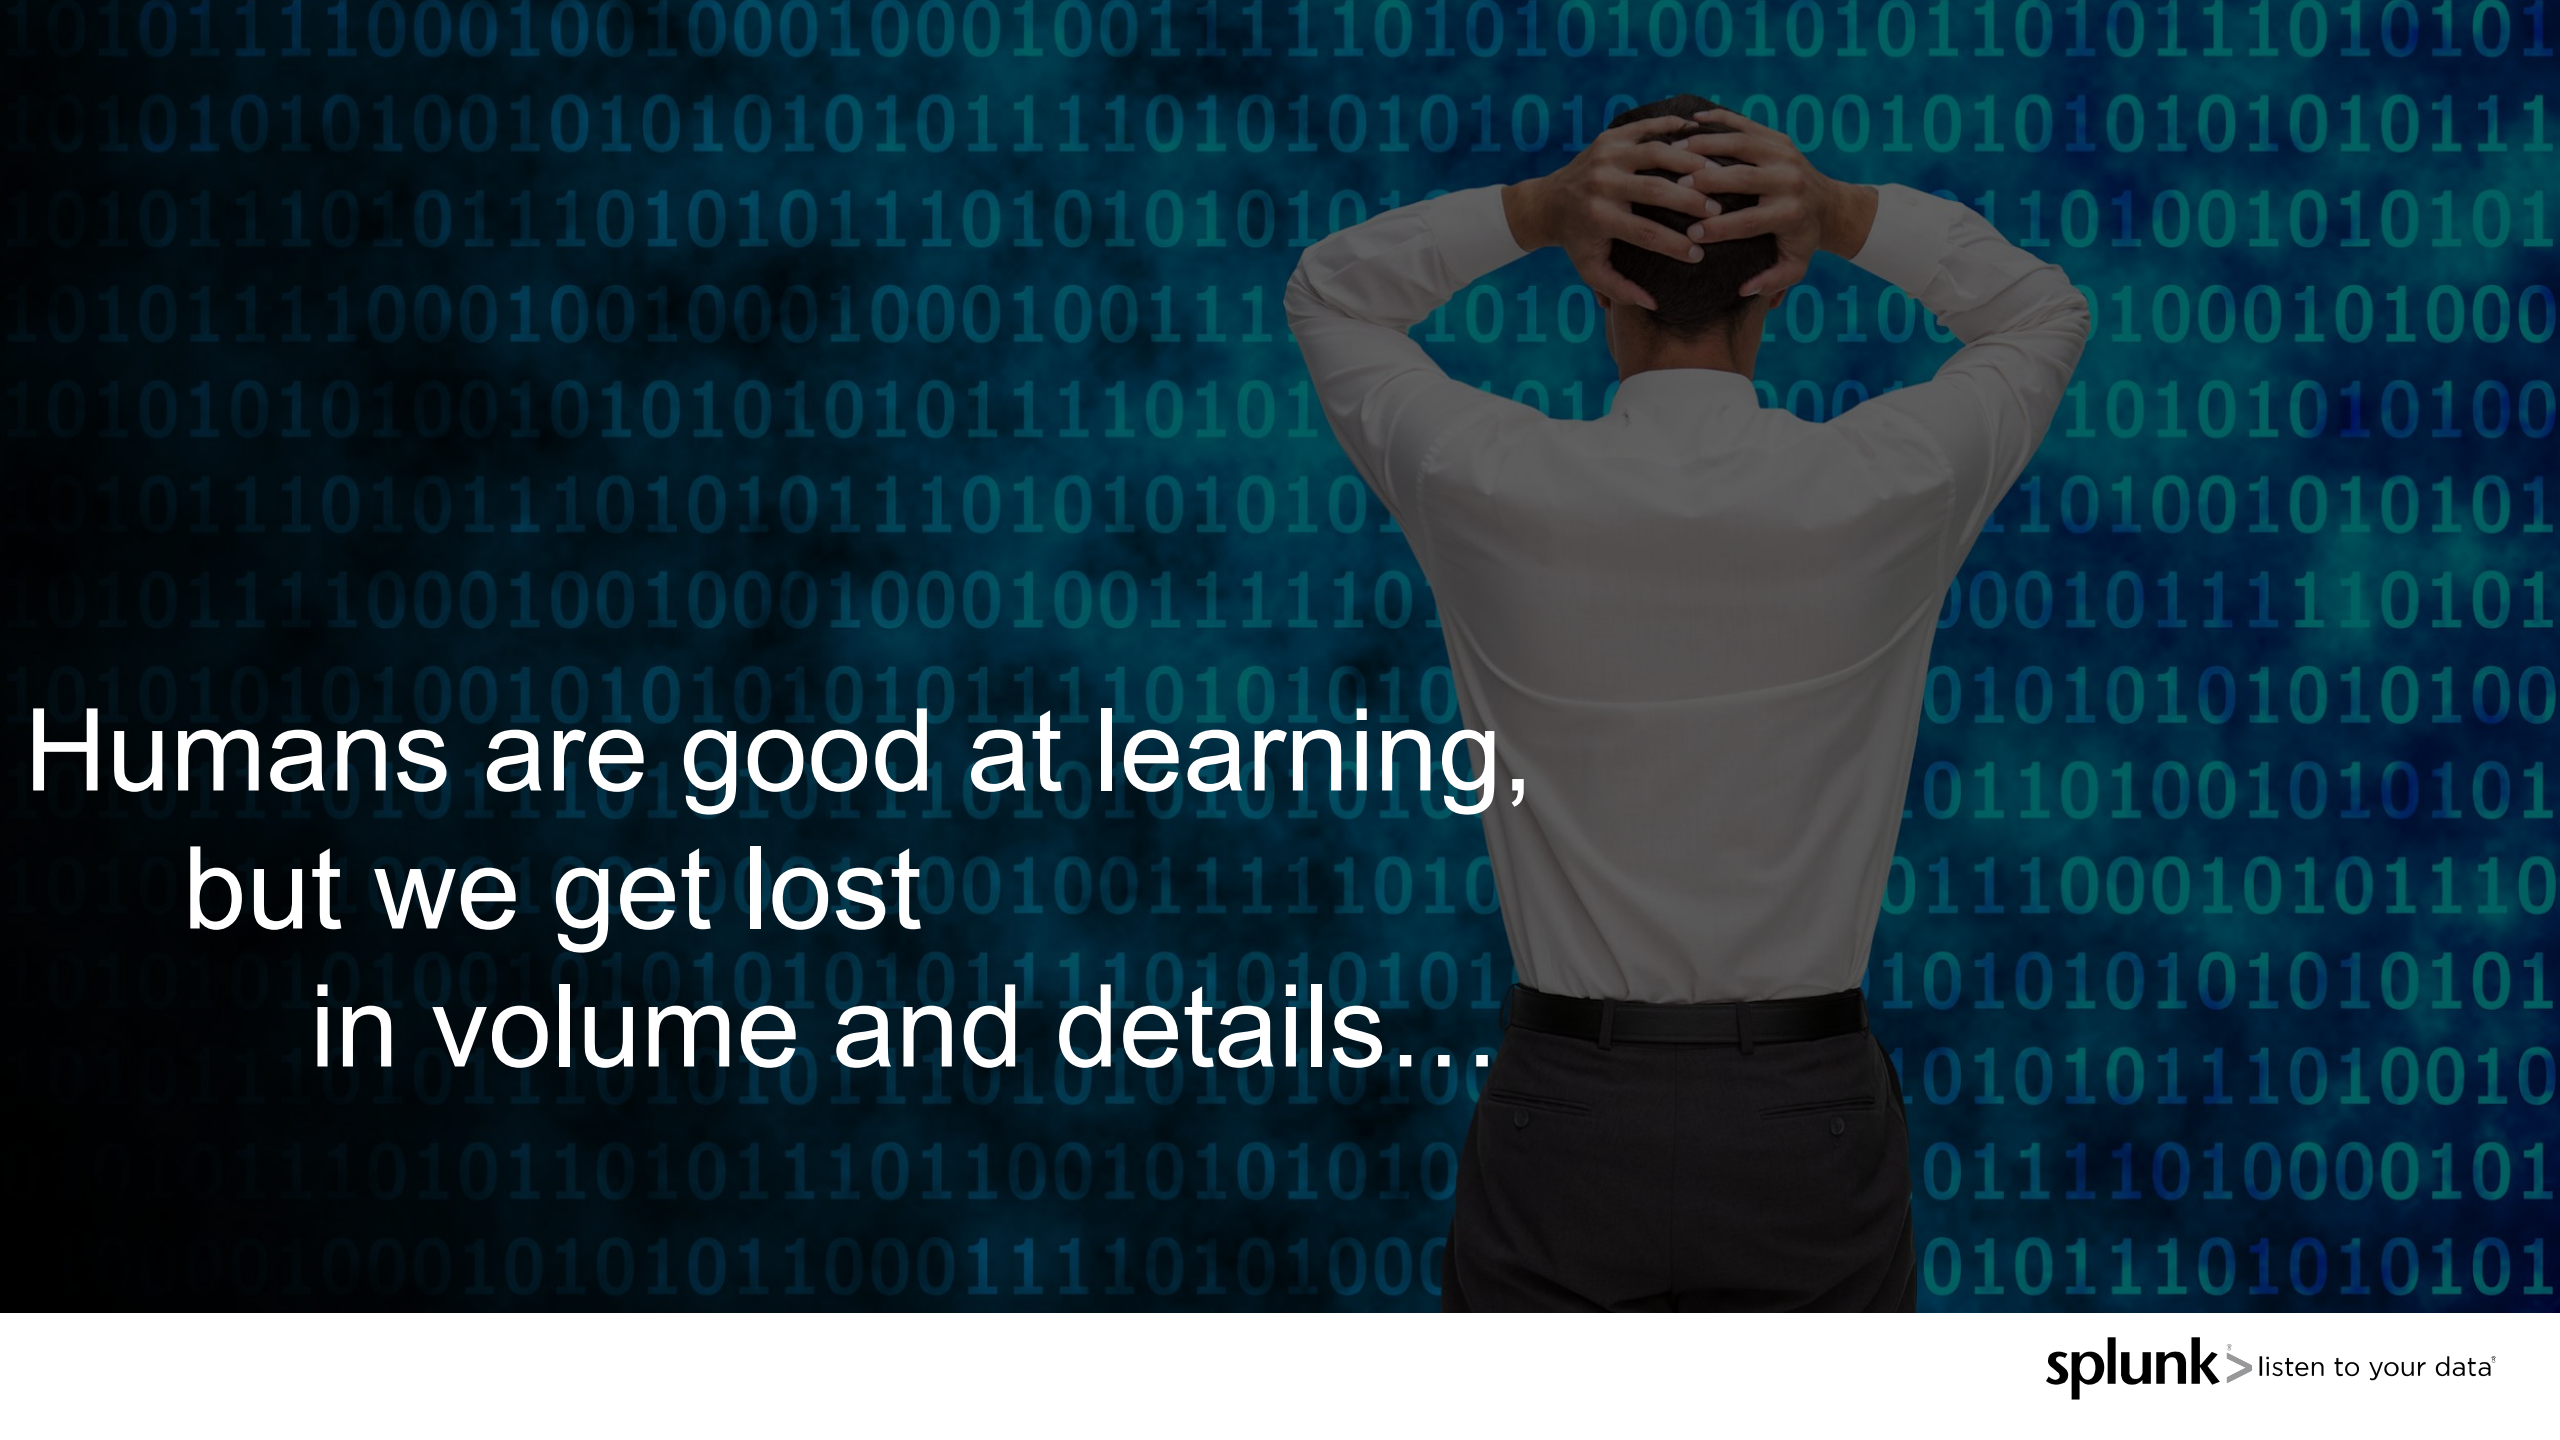


splunk®

What is Splunk?

A T-shirt company that also sells software.



A person in a white shirt is seen from behind, with their hands clasped behind their head. They are standing against a dark blue background filled with a pattern of binary code (0s and 1s) in a lighter blue color. The text "Humans are good at learning, but we get lost in volume and details..." is overlaid on the left side of the image in a white, sans-serif font.

Humans are good at learning,
but we get lost
in volume and details...




What is Machine Learning?

Machine Learning & AI

- ▶ A Function that maps features to an output
- ▶ Learning patterns in your data without being explicitly programmed

Types of ML

- ▶ Supervised
- ▶ Unsupervised
- ▶ Reinforcement



What ML & AI are not

Machine Learning is not Magic

Garbage in = Garbage out

- ▶ Data Scientists spend **80% of their time** cleaning, munging and collecting data
- ▶ Throwing a bunch of data at an algorithm will not result in solving all of your SOC issues
- ▶ Machine Learning requires a solid understanding of statistics and the scientific method

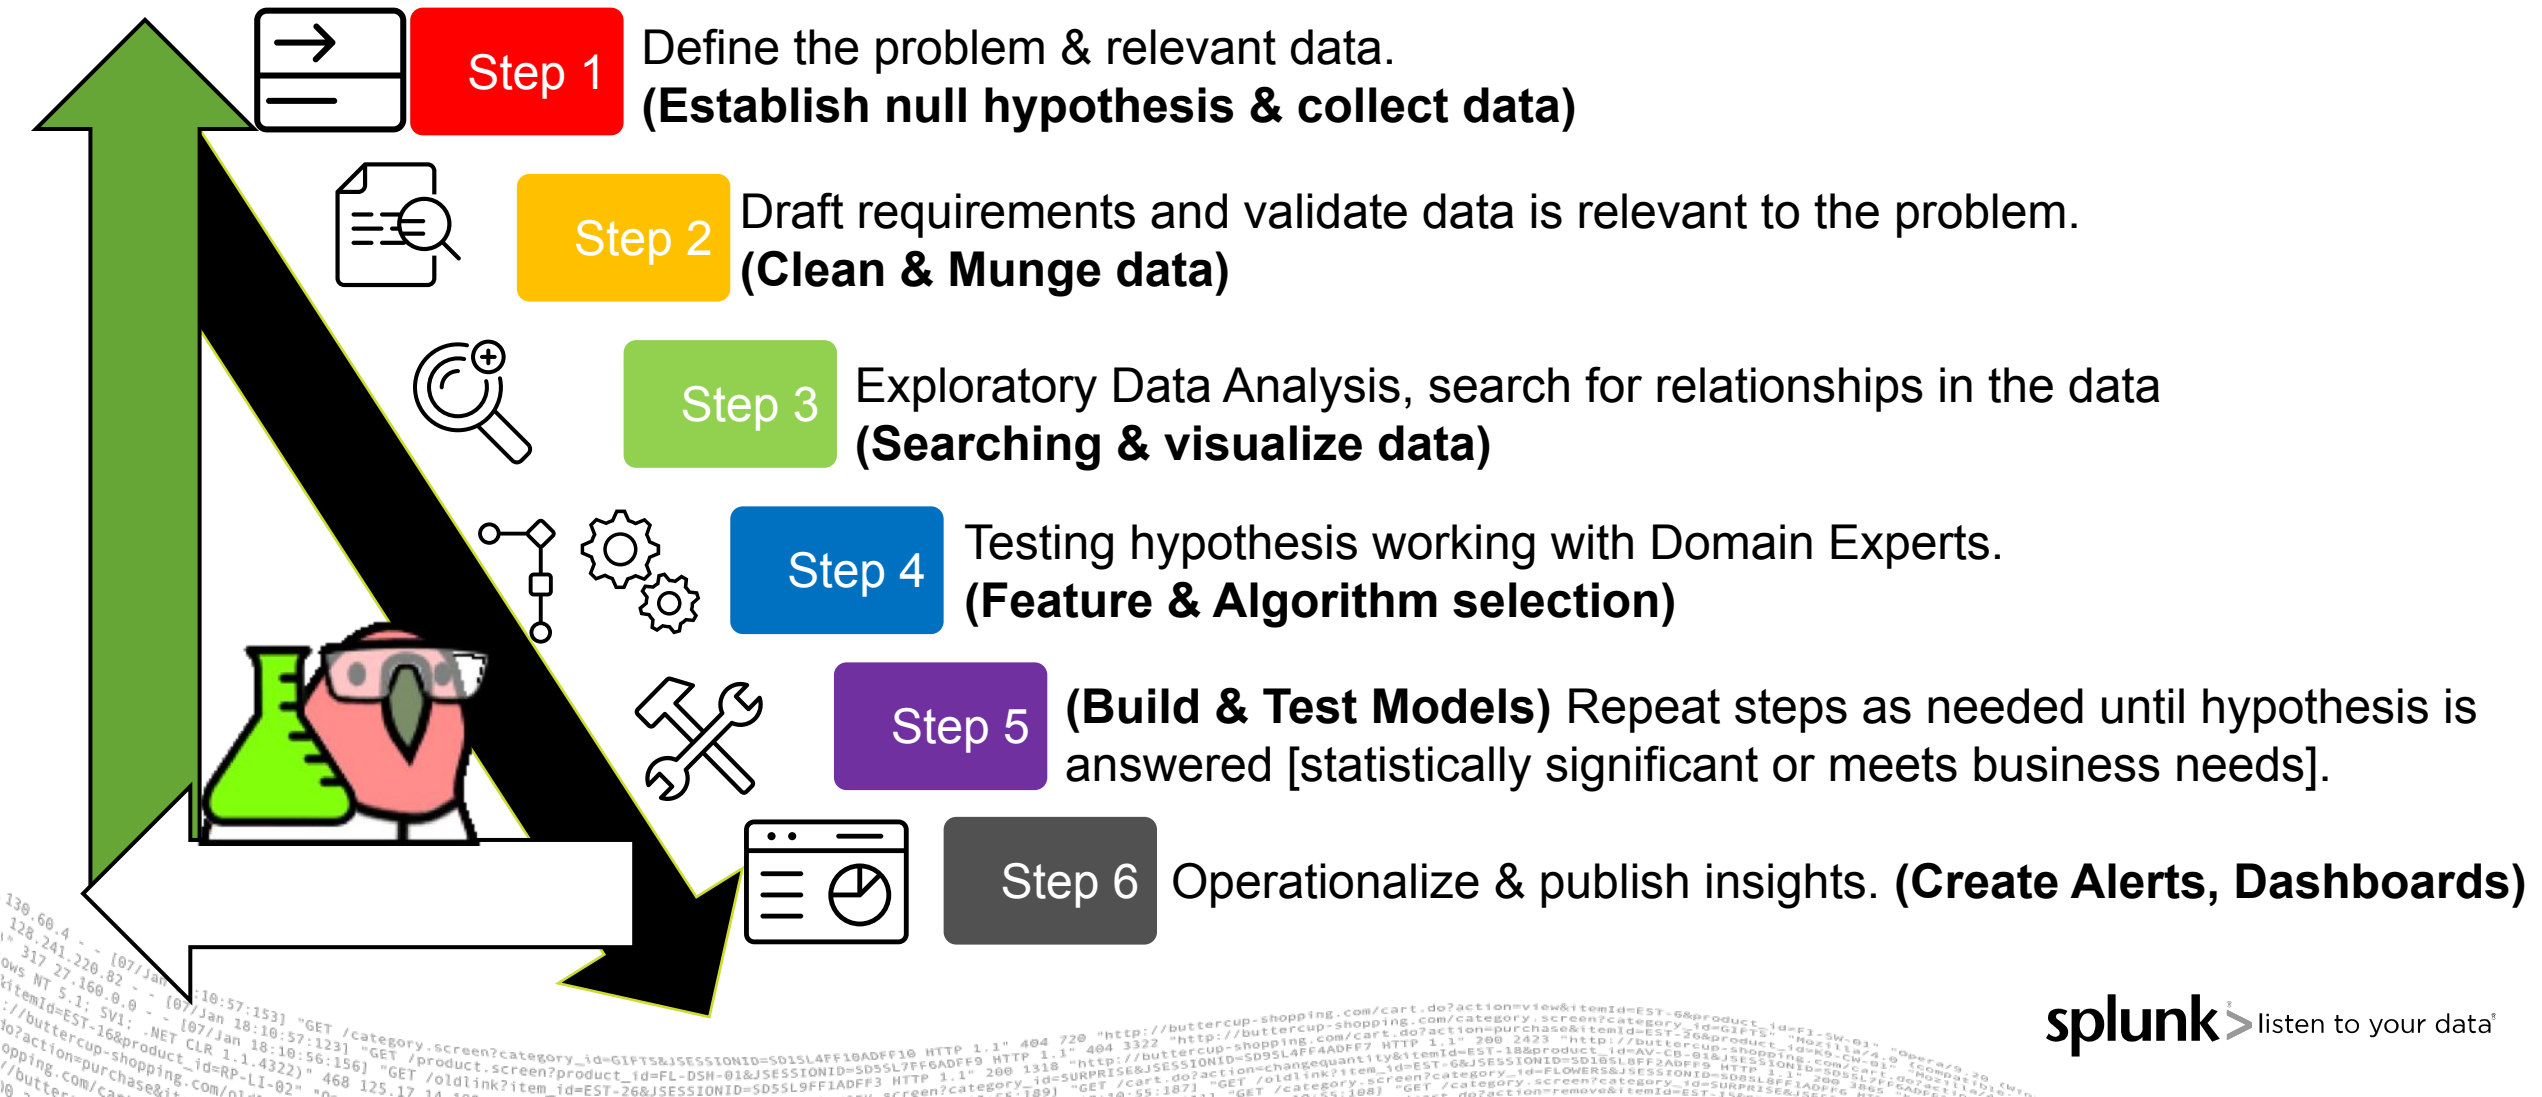
ML & AI require you to understand the fundamental business problem you want to solve.

ML is not a replacement for expert analysts, or engineers. ML requires Subject Matter Experts to enhance security operations and provide feedback to the models.




Data Science Process

What is the problem you're trying to solve?




What can ML do?

Anomaly detection



Deviation from past behavior
Deviation from peers
(aka Multivariate AD or Cohesive AD)
Unusual change in features


Predictive Analytics



Predicting Churn
Predicting Events
Trend Forecasting
Early warning of failure – predictive maintenance

Recommendations (like Netflix)


Clustering



Identify peer groups
Event Correlation
Reduce alert noise
Event Analytics

Use Case: As a NOC/SOC Analyst, I must be alerted when an entity deviates from it's past observed behavior.



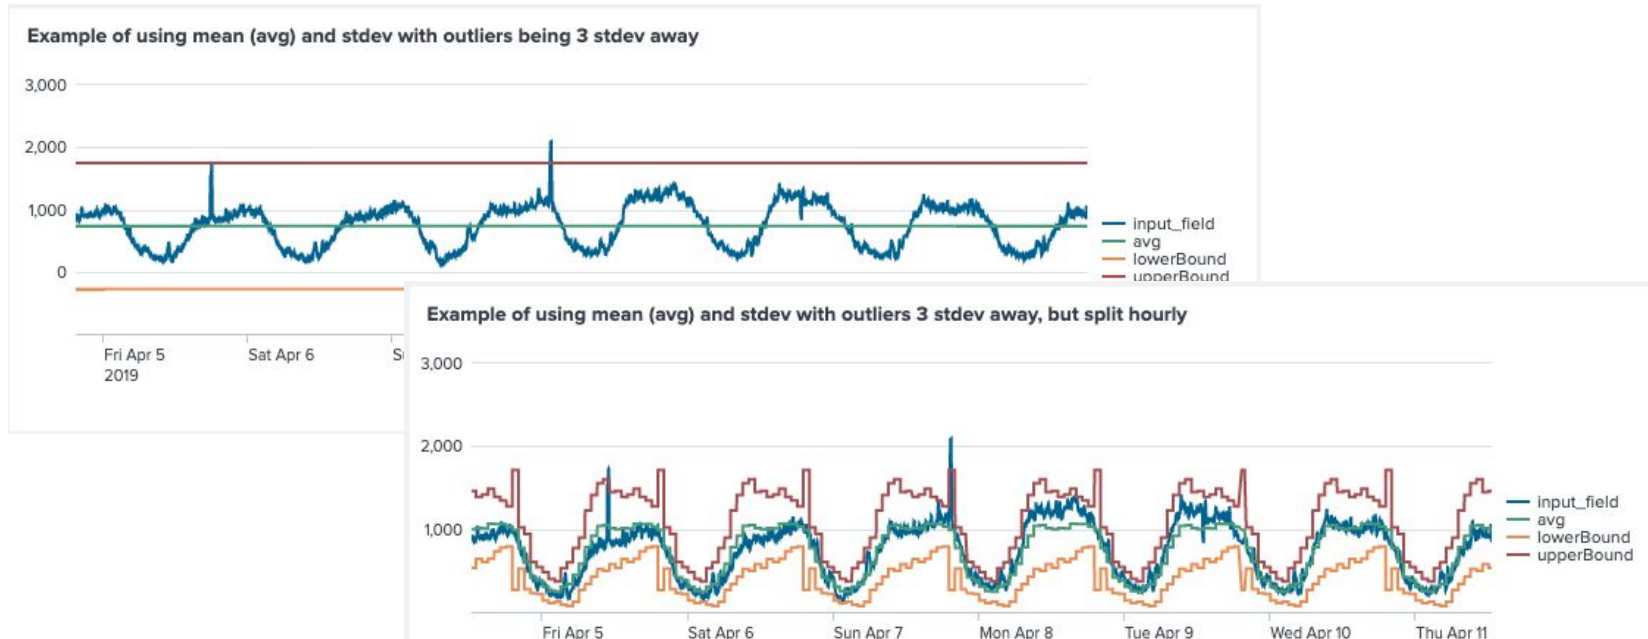
Different behavior on Tuesday Nov 14, returns to normal.



Probability Density Function

Useful Algorithm for determining where numerical outliers will exist.

- ▶ Determine shape of the data: **Normal**, **Exponential**, **Gaussian KDE**
- ▶ Can Understand the difference between **Global Outliers** & **Local Outliers**



Probability Density Function


Use Case Examples:

- ▶ **Account Enumeration/Credential Testing**
 - Abnormally high number of failed logins from device or IP
 - Abnormally high number of account access from device or IP
- ▶ **ATM Transactions / Wire Transfers**
 - Anomalously high number of transactions by merchant
 - Anomalously high transaction by account
- ▶ **Data Exfiltration & Access (Read & Write)**
 - User with high reads & writes to database compared to others in the same role
 - Servers or users with high bytes_out in comparison to peers
- ▶ **IP Theft**
 - High number of requests to API service
 - Speed violations: accounts requesting data at machine speed




What can ML do?

Anomaly detection



Deviation from past behavior
 Deviation from peers
 (aka Multivariate AD or Cohesive AD)
 Unusual change in features


Predictive Analytics



Predicting Churn
 Predicting Events
 Trend Forecasting
 Early warning of failure – predictive maintenance

Recommendations (like Netflix)

Clustering




Identify peer groups
 Event Correlation
 Reduce alert noise

Event Analytics

Predictive Analytics : Proactive

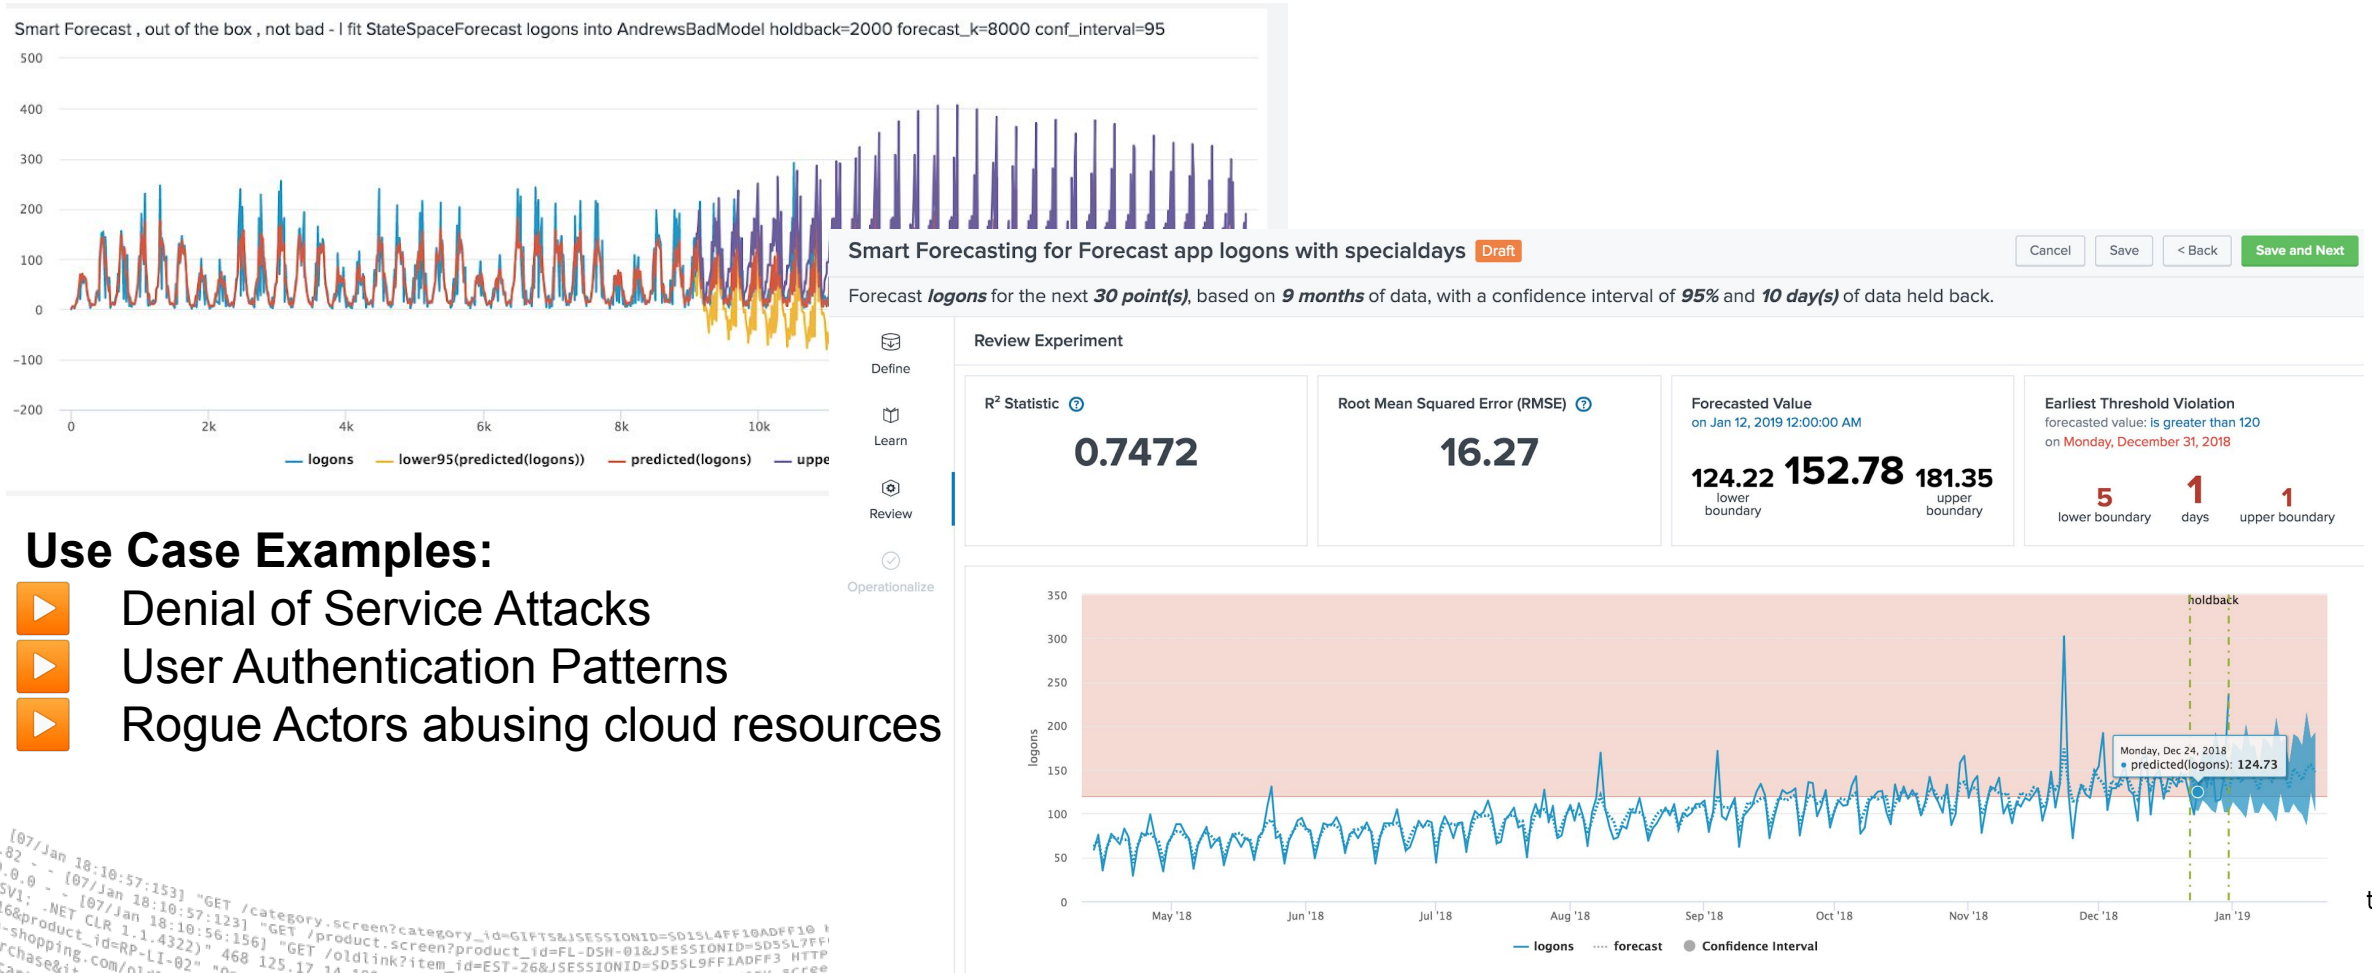
In real time, I update my cloud services usage forecast with the ***meaningful patterns learned from the data***, showing me the next 3 days or so of demand (both the high, low, and actual predicted value).



StateSpace Algorithm


Forecasting learned behaviors that can be leveraged for anomaly detection

Applying forecasting algorithms to security data inform you of trends that are seasonal




What can ML do?

Anomaly detection



Deviation from past behavior
 Deviation from peers
 (aka Multivariate AD or Cohesive AD)
 Unusual change in features


Predictive Analytics



Predicting Churn
 Predicting Events
 Trend Forecasting
 Early warning of failure – predictive maintenance

Recommendations (like Netflix)

Clustering

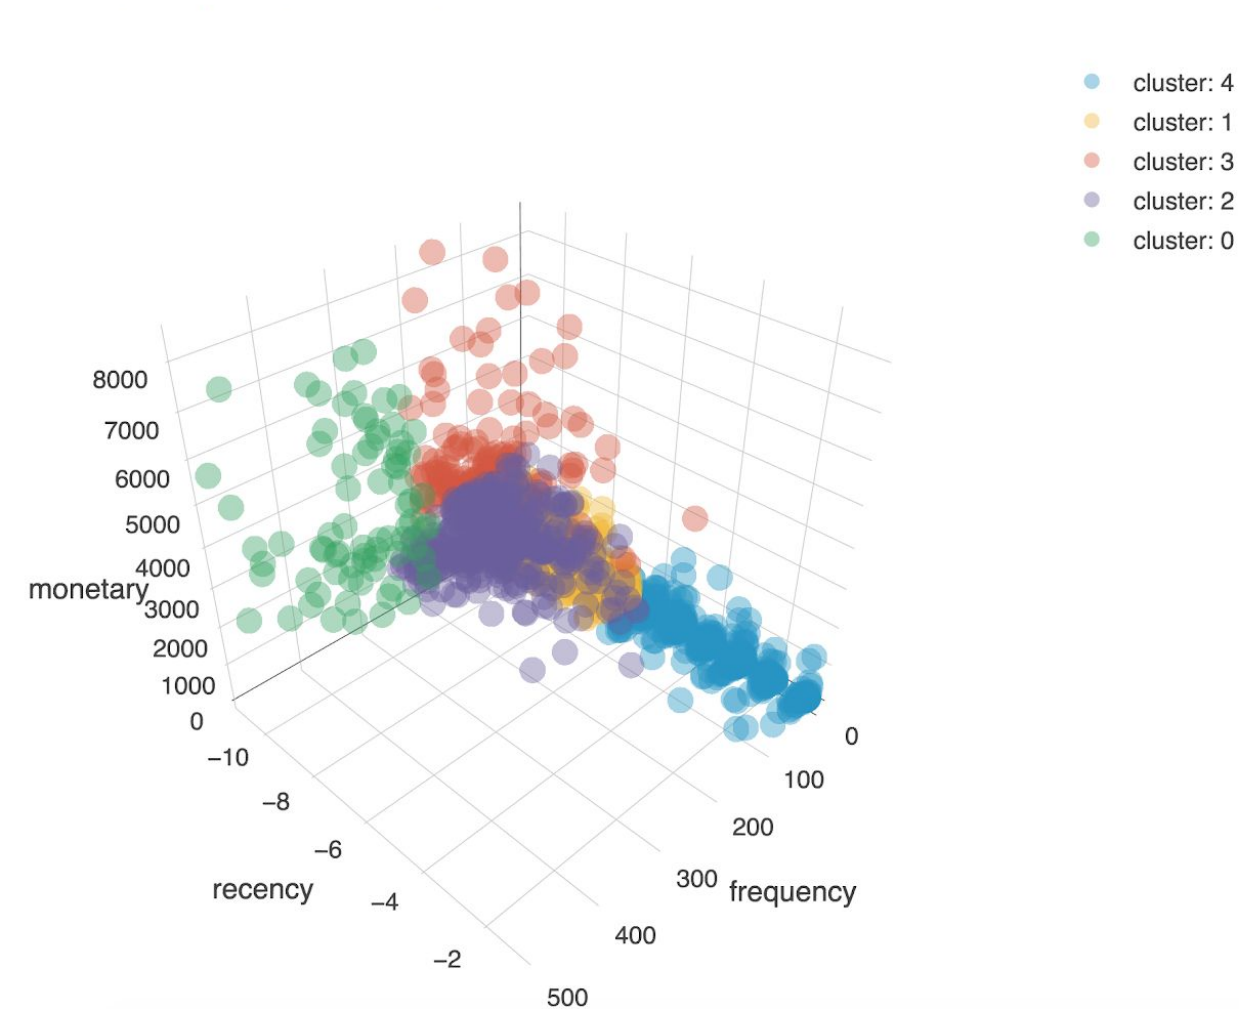


Identify peer groups
 Event Correlation
 Reduce alert noise

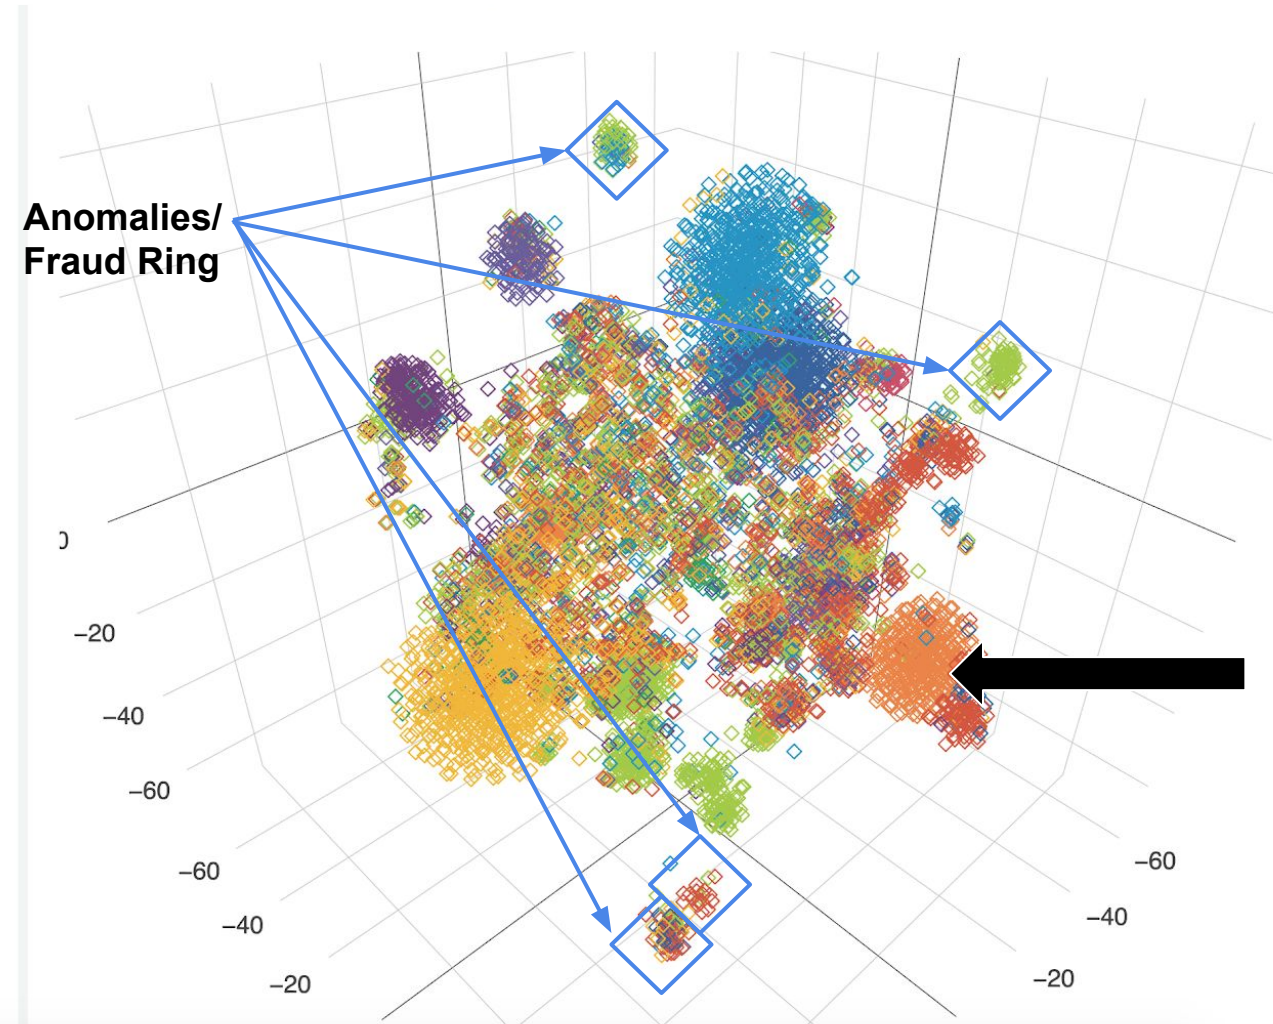
Event Analytics

Clustering : Investigation Outcomes from Data, not Assumptions

Traditional BI - Rule based Clustering



Discovered Behavior Clusters - ML Based Clustering




Clustering Analysis

Objective based look at data to discover where clusters exist based on numerous features.

With an infinite number of dimensions compress the information using Principal Component Analysis to discover where similar groups of entities exist.

Security Use Case - Example features:

- ▶ Payment Card Transactions - Land Speed Violations, Frequency, Recency, Value of Purchases, Duration
- ▶ User Behavior Profiling - AD Group, Systems Accessed, HR data (job function, performance, compensation, etc.)
- ▶ System Profiling - Major OS, Patch Version, User Agent, Services & Software (ssh, sql, rdp, dhcp, etc.)



Exploratory Data Analysis

Use Case Development & Data Science

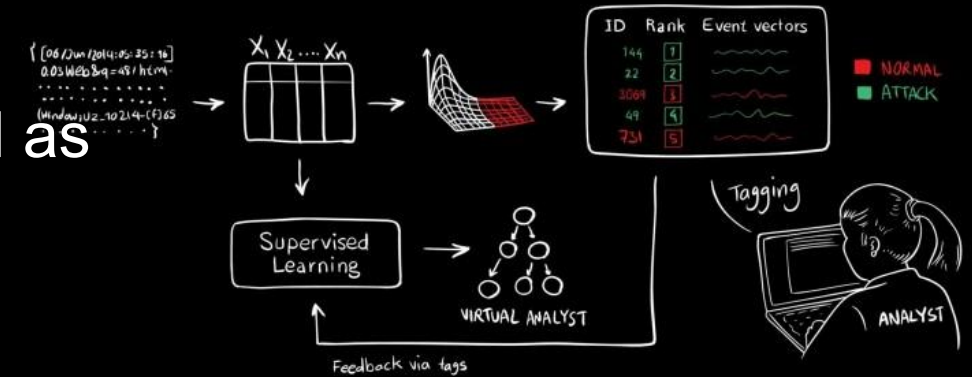
Security Patterns in IT Data

Use Case based approaches to ML/Analytics


What To Look For	Data Source
Abnormally high number of file transfers to USB or CD/DVD	Operating system
Abnormally high number of files or records downloaded from an internal file store or database containing confidential information	File server / Database
Abnormally large amount of data emailed to personal webmail accounts or uploaded to external file hosting site	Email server / web proxy
Unusual physical access attempts (after hours, accessing unauthorized area, etc.)	Physical badge records / Authentication
Excessive printer activity and employee is on an internal watch list as result of demotion / poor review / impending layoff	Printer logs / HR systems
User name of terminated employee accessing internal system	Authentication / HR systems
IT Administrator performing an excessive amount of file deletions on critical servers or password resets on critical applications (rogue IT administrator)	Operating system /Authentication / Asset DB
Employee not taking any vacation time or logging into critical systems while on vacation (concealing fraud)	HR systems / Authentications
Long running sessions, bandwidth imbalance between client & server, Bad SSL Configurations	IPS / IDS / Stream
Known cloud or malware domains, bad SSL Configurations	Threat Intelligence, Custom Lookups
High Entropy Subdomains	Web proxy, DNS, Wiredata

Visualization & Creating Context (EDA)

- ▶ Correlation is used to add context to data I
 - Security issues should not be described as bits, bytes, plaintext or pie charts.
- ▶ Correlation is used to add context to data
 - During EDA or to begin refining your hypothesis.



Destination Traffic- IANA Registration



ps.

ing this activity?"

A world map with green lines connecting three points: one in North America, one in Europe, and one in East Asia. The lines represent connections between these three locations.

- ### Attempted SSH Access by Country



Number of Connections: 47616

SSH Attempts - Numerical Outliers

Machine Learning & Security Analytics

DGA Domains

Domain Generating Algorithms (DGA)

What's DGA?

“A Domain Generating Algorithm (DGA) is a program or subroutine that provides malware with new domains on demand or on the fly.”

Challenges to detect DGAs:

- ▶ Static matching runs against potentially infinite blacklist entries $O(\infty)$
- ▶ Regex can narrow down this list, but still hard to compute and find rules (and define exceptions for rules)
- ▶ Unknown unknowns?
- ▶ Want to get fuzzy?
- ▶ Good use case for Statistics/ML!

▶ Example of DGAs:

domain ↕

iuqerfsodp9ifjaposdfjhgosurijfaewrrergwea

ifferfsodp9ifjaposdfjhgosurijfaewrrergwea

ayylmaotjhsstasdfsdfasdfsdfasdfsdfasdfsdf

lazarusse.suiche.sdfjhgosurijfaqwqwqrgwea

sdfjhgosurijfaqwqwqrgwea

Example IoCs for Wannacry

(<https://cert.europa.eu/static/SecurityAdvisories/2017/CERT-EU-SA2017-012.pdf>)

- **Botnet sends queries with 16 letters randomly prepended to the victim's domain.**
 - xyuiasdfcosic.www.halpme.com
 - alkdfejenjasd.www.halpme.com

A pixelated red alien character with a large, rounded body, small eyes, and a wide, toothy grin. It has a small, pointed head and a single leg visible. The character is rendered in a simple, blocky style with a limited color palette.

- **Advanced malware uses a Domain Generation Algorithm (Random Subdomain)**
 - d0290d00xasdf.no-ip[.]org

- **DNS Tunneling (Query)**
 - **dnscat.912701a98e9bde415c4ad70007beaf54d2**
 - **dnscat.925401a98ebe0cf540b20d001a4b5e726494b001bb4c192bb68fe73c000bf7c1c0e**

- **Shannon Entropy of DNS Query or HTTP destination**
- **Character Length of DNS Query or HTTP destination**

- “... a measure of uncertainty in a random variable”

- The more random a string is, the higher its calculation of randomness.

- *aaaaaa.com* (Score 1.8)
- *Google.com* (Score 2.6)
- *lc49f66b73141b5c1.com* (Score 4.1)
- Domains and subdomains with high entropy are good indicators of malicious behavior.
- **We can filter to domains or subdomains with a score above 3 or 4.**

$$H = - \sum p(x) \log p(x)$$




Text Mining Approach

n-gram distribution + principal component analysis

n-gram model is a type of probabilistic language model for predicting the next item in such a sequence in the form of a $(n - 1)$ -order Markov model.

English Bigrams			Domain Bigrams	
LETTER	FREQUENCY		LETTER	FREQUENCY
th	0.03883		in	0.01702
he	0.03681		er	0.01550
in	0.02284		an	0.01333
er	0.02178		re	0.01290
an	0.02141		es	0.01271
re	0.01749		ar	0.01188
nd	0.01572		on	0.01135
on	0.01418		or	0.01051
en	0.01383		te	0.01017
at	0.01336		al	0.00976
ou	0.01286		st	0.00921
ed	0.01276		ne	0.00921
ha	0.01275		en	0.00897



<https://en.wikipedia.org/wiki/N-gram>

<https://www.semanticscholar.org/paper/Detecting-DNS-Tunnels-Using-Character-Frequency-Born-Gustafson/c7cc7c16e8952facae1e4dfb0dd768a4504cd5cb>

Feature Engineering

domain	class	subclass	ut_consonant_ratio	ut_digit_ratio	ut_domain_length	ut_meaning_ratio	ut_shannon	ut_vowel_ratio	PC_1	PC_2	PC_3
lvmaehe1voogfbss.net	dga	chinad	0.600	0.050	20.000	0.300	3.784	0.300	0.502	-0.304	0.092
1amnl1a519ort3p12o09111e6288k.com	dga	newgoz	0.281	0.531	32.000	0.156	3.925	0.188	-0.358	-0.008	0.181
fiabg19j4wxu16sacop1su49dx.org	dga	newgoz	0.516	0.258	31.000	0.226	4.196	0.226	0.102	0.763	0.415
fspfffyddxni.pl	dga	locky	0.900	0.000	15.000	0.067	3.107	0.067	0.044	0.072	-0.066
ulpkn41fwor3pyqv9551j4f35c.com	dga	newgoz	0.600	0.333	30.000	0.067	4.282	0.100	-0.362	-0.001	0.177
aaqa93u5uybd1nbe.net	dga	chinad	0.500	0.200	20.000	0.300	3.684	0.350	0.659	-0.385	0.117
hao6m700qnr07d3y.cn	dga	chinad	0.526	0.316	19.000	0.105	3.827	0.158	-0.063	0.013	-0.019
1y1j69jb62wpg1h58kdp3mb8n2.org	dga	newgoz	0.600	0.400	30.000	0.067	4.282	0.033	0.178	0.823	0.404
play.googleapis.com	legit	legit	0.600	0.000	19.000	0.579	3.471	0.368	-0.222	-0.085	0.037
051i8937btzxhotb.info	dga	chinad	0.476	0.333	21.000	0.286	4.011	0.190	0.049	0.117	-0.125

« prev 1 2 3 4 5 6 7 8 9 10 next »

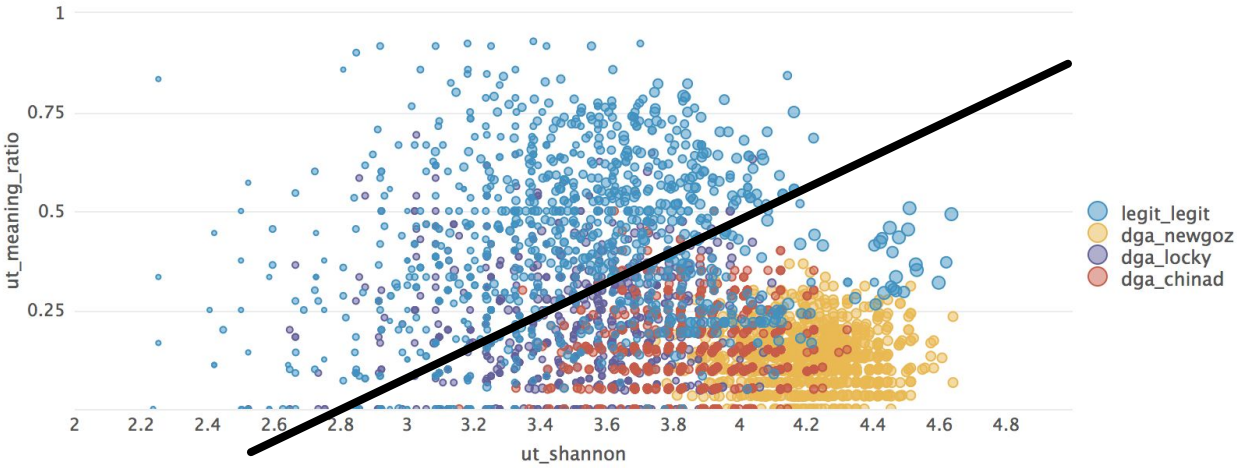
▶ More text based features can significantly improve your machine learning models

- Be wary of overfitting!


▶ Examples feature engineering ideas

- (e.g. count of subdomains, age of domain registration, rating/scoring from threatlists for known malicious domains etc.)

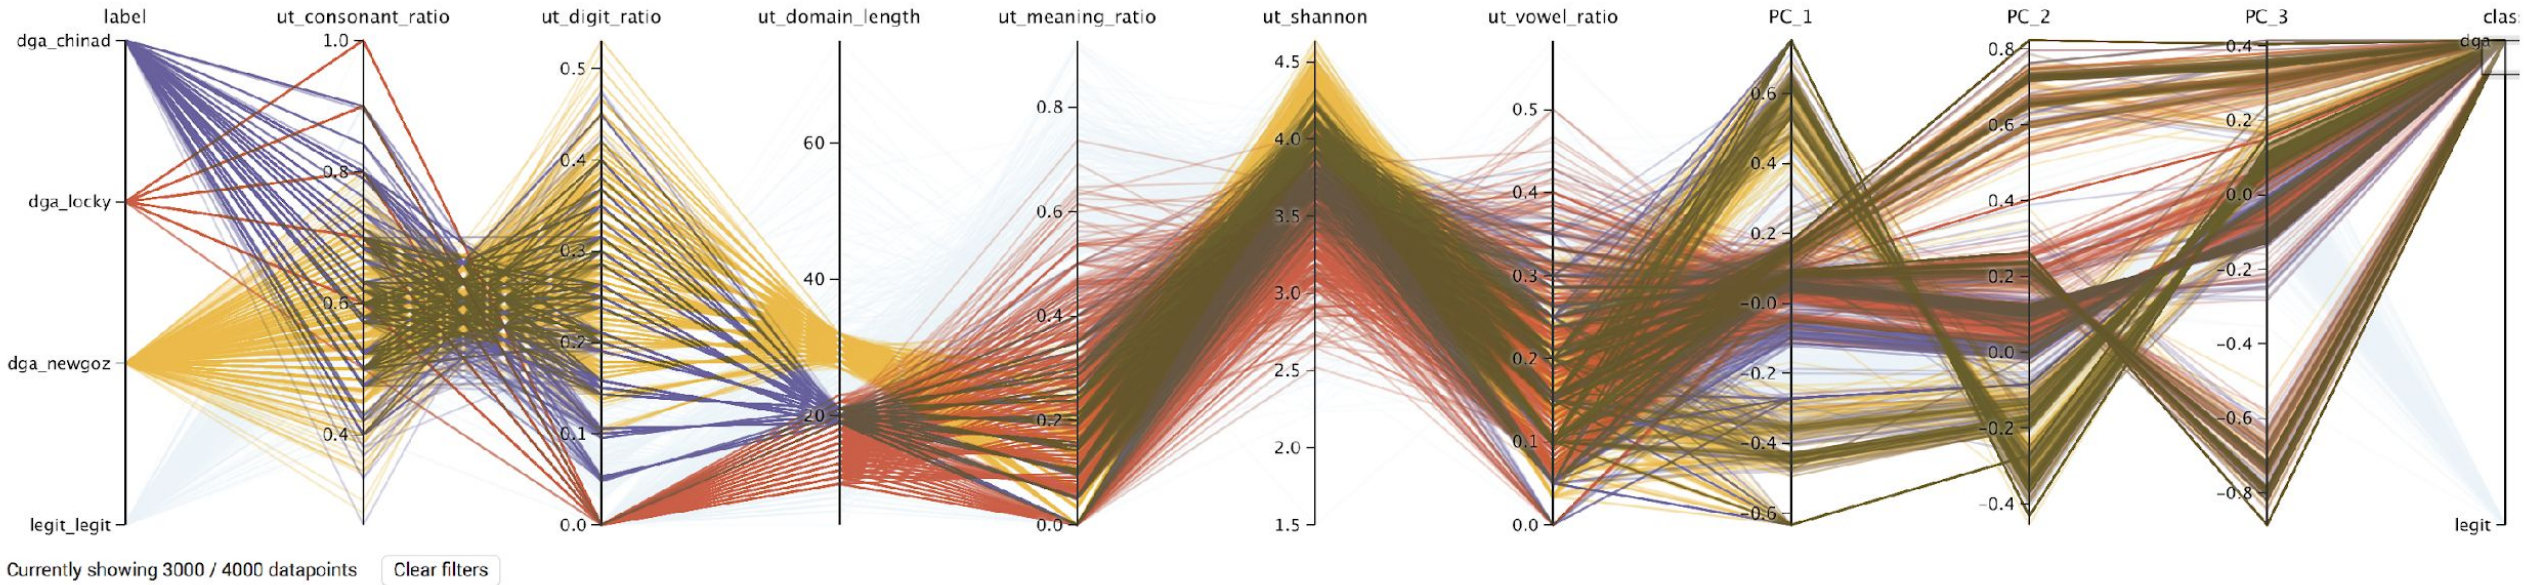
Distribution of classes depending on example feature combination



Identify useful features for classification with the analyzefields command




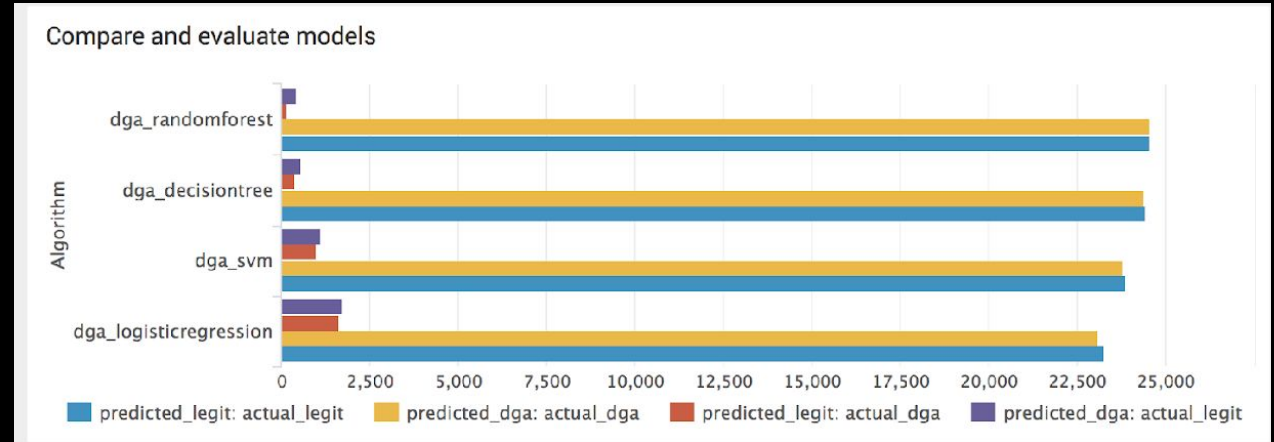
Parallel coordinate chart of classes and top features



Training & Testing Models

Selecting the right algo

- ▶ Goal of machine learning is to enhance security operations not add to the “alert fatigue”.
- ▶ Some Algorithms will be better than others for certain types of problems
- ▶ Minimize False Positives
- ▶ Accept/Reduce risk(s) associated with False Negatives



Operationalize

Use the model against new data!

Count of predictions

6




Trend of true predictions

6  -8

Trend of false predictions

0  -1

Prediction performance over time



Results of machine learning algorithm (dga_randomforest) applied to new domains

_time	domain	class	predicted(class)
2017-09-19 22:10:59	mrop8i1scak54s2hskfpq443z.org	dga	dga
2017-09-19 22:10:58	stackadapt.com	legit	legit
2017-09-19 22:10:57	cf.dropboxstatic.com	legit	legit
2017-09-19 22:10:54	8fjsfjlapjxhcm.org	dga	dga
2017-09-19 22:10:52	g5te08189ly791v20rrr53dsdv.org	dga	dga

Results of machine learning algorithm (dga_randomforest) with DGA detected

_time	domain	class
2017-09-19 22:10:59	mrop8i1scak54s2hskfpq443z.org	dga
2017-09-19 22:10:54	8fjsfjlapjxhcm.org	dga
2017-09-19 22:10:52	g5te08189ly791v20rrr53dsdv.org	dga
2017-09-19 22:10:44	ljydrwun.info	dga
2017-09-19 22:10:43	txteviaggwqudpzh.info	dga

Improve Model

Human in the Loop Feedback

- ▶ Not every alert or prediction is going to be correct.
- ▶ This is where SMEs are needed to provide feedback to the model for retraining to increase the accuracy.
- ▶ SMEs can also help engineer new features as malware evolves and tactics change

	time	datetime	class	domain	key_domain
1	1505852183.609000	09/19/17 22:16:23	legit	dimmhjx.xyz	LEGIT DGA
2	1505852180.980000	09/19/17 22:16:20	legit	rvxoudaurvjf.info	LEGIT DGA
3	1505852159.000000	09/19/17 22:15:59	dga	1qntxfsl3eloj8hdhbkdk1qdcfqf.org	LEGIT DGA
4	1505852157.000000	09/19/17 22:15:57	dga	nucv6amdxse1vbtu.biz	LEGIT DGA
5	1505852154.000000	09/19/17 22:15:54	dga	qeuextlwmjg.info	LEGIT DGA
6	1505852147.000000	09/19/17 22:15:47	dga	14fb5x4pu2zmu12eulks162u7b3.com	LEGIT DGA
7	1505852137.000000	09/19/17 22:15:37	dga	f3upm510ybndfqycfcz1ajbghu.org	LEGIT DGA
8	1505852136.000000	09/19/17 22:15:36	dga	lv31si318e57gk1gdcsi114t5m9.com	LEGIT DGA
9	1505852134.000000	09/19/17 22:15:34	dga	un905fm8cfb9etmx23m8sy5y.net	LEGIT DGA
10	1505852132.000000	09/19/17 22:15:32	dga	m3e3ytfvqgtj1wv1d3ka0zf3j.net	LEGIT DGA

« prev 1 2 3 4 5 6 7 8 9 10 next »

Reality check: Detect Unknown Unknowns?

Example WannaCry

- ▶ Check how our trained model performs against WannaCry C&C domains that the model has **NOT** been trained on.
- ▶ Model predictions can be made actionable immediately with Splunk Alerts or turn into notable event frameworks like in ES



Check if Wannacry would have been detected

Source of IoC (domain names): <https://cert.europa.eu/static/SecurityAdvisories/2017/CERT-EU-SA2017-012.pdf>

_time	domain	class	predicted(class)
2017-06-01 23:13:57	iuqerfsodp9ifjaposdfjhgosurijfaewrwergwea	unknown	dga
2017-06-01 23:13:57	ifferfsodp9ifjaposdfjhgosurijfaewrwergwea	unknown	dga
2017-06-01 23:13:57	ayylmaotjhsstasdfasdfasdfasdfasdfasdf	unknown	dga
2017-06-01 23:13:57	lazarusse.suiche.sdfjhgosurijfaqwqwrgwea	unknown	dga
2017-06-01 23:13:57	sdfjhgosurijfaqwqwrgwea	unknown	dga

Splunk CTF – July 25th 2019

Brisbane | Sydney | Melbourne | Canberra | Adelaide | Perth | Hobart | Darwin | Auckland | Wellington



Australia & New Zealand Boss of the SOC Day 2019



A ntipodean Splunkers! Ever the trailblazers, Australia and New Zealand are full steam ahead for our [Australia & New Zealand Boss of the SOC \(BOTS\) Day](#) held on July 25, 2019. Splunk Recently held its North American BOTS day in June, with hundreds of participants right across the country. Technically [we did it first](#), but we're happy to give our North American family credit (for now).



What is Boss of the SOC?

```
Opera/9.01 (Windows NT 5.1; U; en)" 539 10.2.1.44 [07/Jun 18:10:42:109] "GET 12.130.60.4 [07/Jun 18:10:57:153] "GET /category
on=view&itemId=EST-6 product_id=FI-SW-01" "Opera/9.20 (Windows NT 6.0; U; en)" 559 128.241.220.82 [07/Jun 18:10:57:123] "GET /product
ry.screen?category_id=GIFTS" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1)" 317 27.160.0.0 [07/Jun 18:10:56:156] "GET /oldlink?item
ase&itemId=EST-26 product_id=K9-CW-01" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; SV1; .NET CLR 1.1.4322)" 468 125.17.14.100
http://buttercup-shopping.com/cart.do?action=addtocart&itemId=EST-16&product_id=RP-LI-02" "Opera/9.20 (Windows NT 6.0; U; en)"
ONID SD5SL7FF6ADFF10 HTTP 1.1" 404 3860 "http://buttercup-shopping.com/oldlink?item_id=EST-18" "Opera/9.20 (Windows NT 6.0; U; en)" 766 130.253.37.97
```

Thank You





Install ML Toolkit

Step 1: install ML Toolkit app

<https://splunkbase.splunk.com/app/2890>




Step 2: install Python for Scientific Computing add-on

<https://splunkbase.splunk.com/app/2890/#/details>

Step 3: restart Splunk

That's it! Explore the ML Toolkit:

Prediction, Outlier Detection, Forecasting, Clustering
 Showcase examples for IT/Security/Business/IoT use cases
 Assistants: use your own data, build models, view in search



Resources

who I owe credit to:

- ▶ Philipp Drieger: DGA App & Content <https://splunkbase.splunk.com/app/3559/>
 - Conf17 Presentation Recording – <http://conf.splunk.com/files/2017/recordings/automating-threat-hunting-with-machine-learning.mp4>
- ▶ Mike Fisher: Building a crystal ball
 - Conf16 Presentation – <https://conf.splunk.com/files/2016/slides/building-a-crystal-ball-forecasting-future-values-for-multi-cyclic-time-series-metrics-in-splunk.pdf>
- ▶ Macy Cronkite: Anomaly Hunting with Splunk
 - Conf16 Presentation – <https://conf.splunk.com/files/2016/slides/anomaly-hunting-with-splunk-software.pdf>
- ▶ Xander Johnson & Zidong Yang: ML API
 - Conf17 Presentation – <http://conf.splunk.com/files/2017/slides/advanced-machine-learning-using-the-extensible-ml-api.pdf>

- ▶ Andrew Stein
 - General ML advice & mentoring

40

- splunk**  listen to your data™