Death to the IOC

What's next in Threat Intelligence

Bhavna Soman, Microsoft Defender Research

whoami

- Microsoft Defender Research
- Past: Threat Intelligence and APT response
- Present: Builds algorithms to classify malware in real-time
- Future: ??

Threat Analysis is largely manual work

- Not just automated IOC feeds
- Some Reverse Engineering, Forensics
- Lots of… READING!!

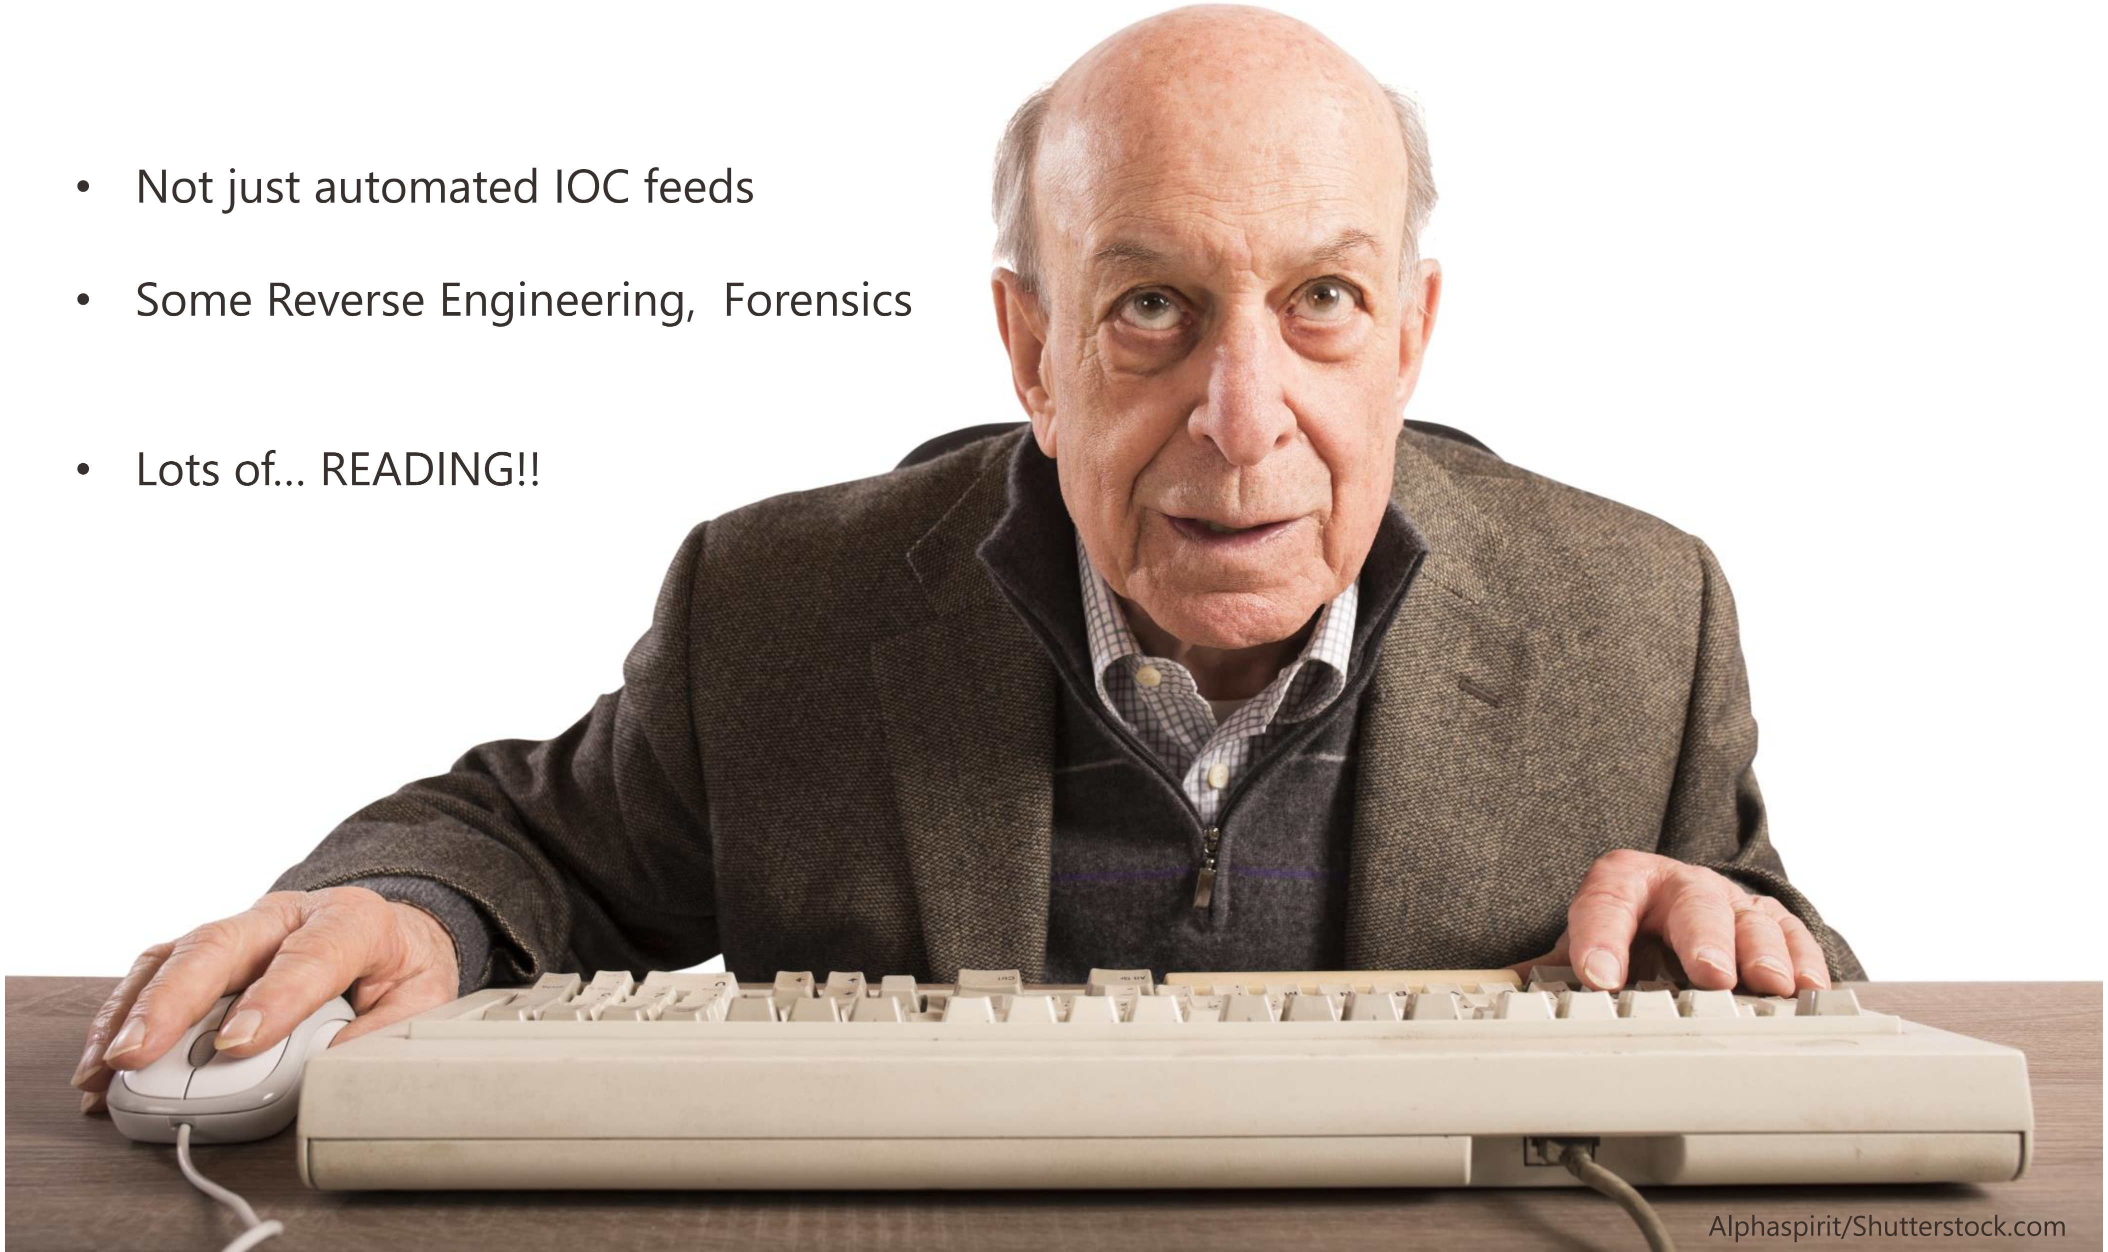

Alphaspirit/Shutterstock.com

KEY FINDINGS

APT28 targets insider information related to governments, militaries, and security organizations that would likely benefit the Russian government.

GEORGIA	EASTERN EUROPE	SECURITY ORGANIZATIONS
APT28 likely seeks to collect intelligence about Georgia's security and political dynamics by targeting officials working	APT28 has demonstrated interest in Eastern European governments and security organizations. These victims	APT28 appeared to target individuals affiliated with European security organizations and global multilateral

https://www.fireeye.com/content/dam/fireeye-www/global/en/current-threats/pdfs/rpt-apt28.pdf

Dragonfly: Western energy sector targeted by sophisticated attack group

Resurgence in energy sector attacks, with the potential for sabotage, linked to re-emergence of Dragonfly cyber espionage group.

https://www.symantec.com/blogs/threat-intelligence/dragonfly-energy-sector-cyber-attacks

NEWS

Russian Cozy Bear APT 29 hackers may be impersonating State Department

Russian Cozy Bear hackers may be impersonating the U.S. State Department in a large, new spear-phishing campaign, plus other cybersecurity news.

https://www.csoonline.com/article/3321911/security/russian-cozy-bear-apt-29-hackers-may-be-impersonating-state-department.html

Shoebox of Tools, Tactics and Procedures

APT 28

Data Obfuscation, Connection Proxy, Standard Application Layer Protocol, Remote File Copy, Rundll32 ,Indicator Removal on Host, Timestomp, Credential Dumping, Screen Capture, Bootkit, Component Object Model Hijacking, Exploitation for Privilege Escalation, Obfuscated Files or Information, Input Capture, Replication Through Removable Media, Communication Through Removable Media, Pass the Hash, Data Staged, Data from Removable Media, Peripheral Device Discovery, Access Token Manipulation, Valid Accounts, Office Application Startup, System Owner/User Discovery, Process Discovery, System Information Discovery, File Deletion, Credentials in Files, File and Directory Discovery, Network Sniffing, Dynamic Data Exchange, Data from Local System, Hidden Files and Directories, Scripting, Logon Scripts, Spearphishing Attachment, Deobfuscate/Decode Files or Information, Exploitation of Remote Services, Exploitation for Defense Evasion, Data from Information Repositories, User Execution

APT 29

PowerShell, Scripting, Indicator Removal on Host, Software Packing, Scheduled Task, Registry Run Keys / Start Folder, Bypass User Account Control, Windows Management Instrumentation Event Subscription, Windows Management Instrumentation, Pass the Hash, Accessibility Features, Domain Fronting ,Multi-hop Proxy, Spearphishing Attachment, Spearphishing Link, Exploitation for Client Execution, User Execution

Dragonfly

Screen Capture, PowerShell, Remote File Copy, File Deletion, Create Account, Disabling Security Tools, External Remote Services, Brute Force, Credential Dumping, Scripting, Masquerading, Indicator Removal on Host, Web Shell, Commonly Used Port, Email Collection, Remote Desktop Protocol, Network Share Discovery, Scheduled Task, Forced Authentication

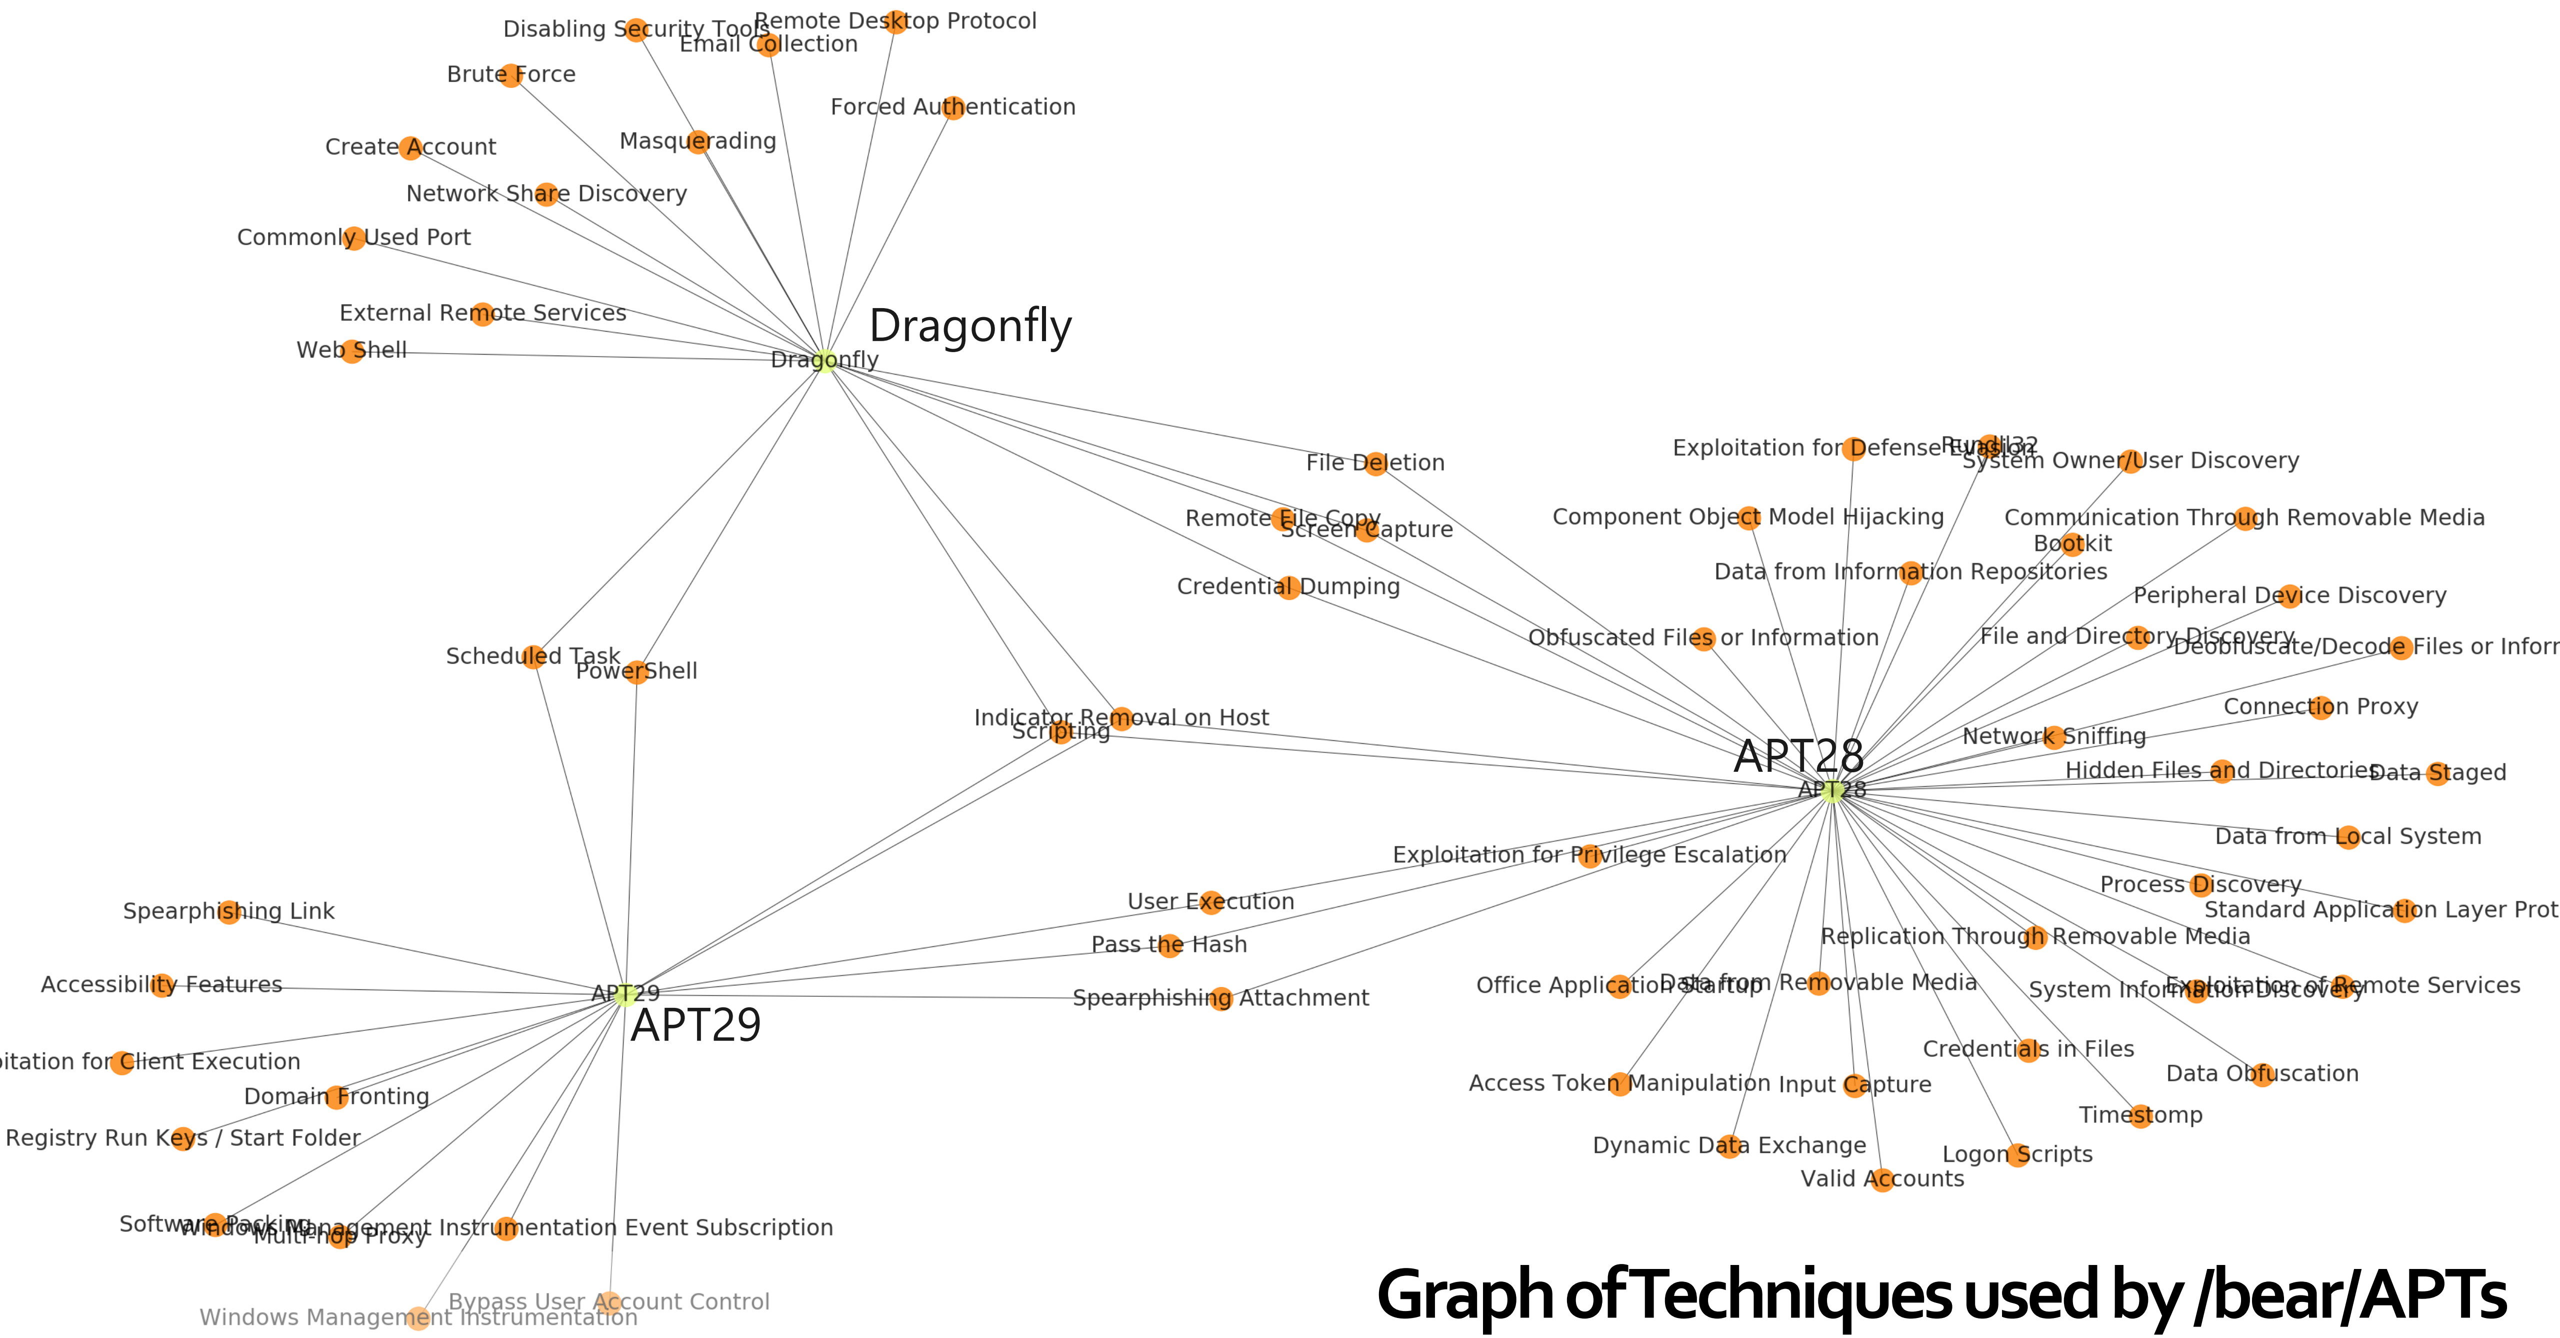

Graph of Techniques used by /bear/APTs

https://mitre-attack.github.io/caret/#/

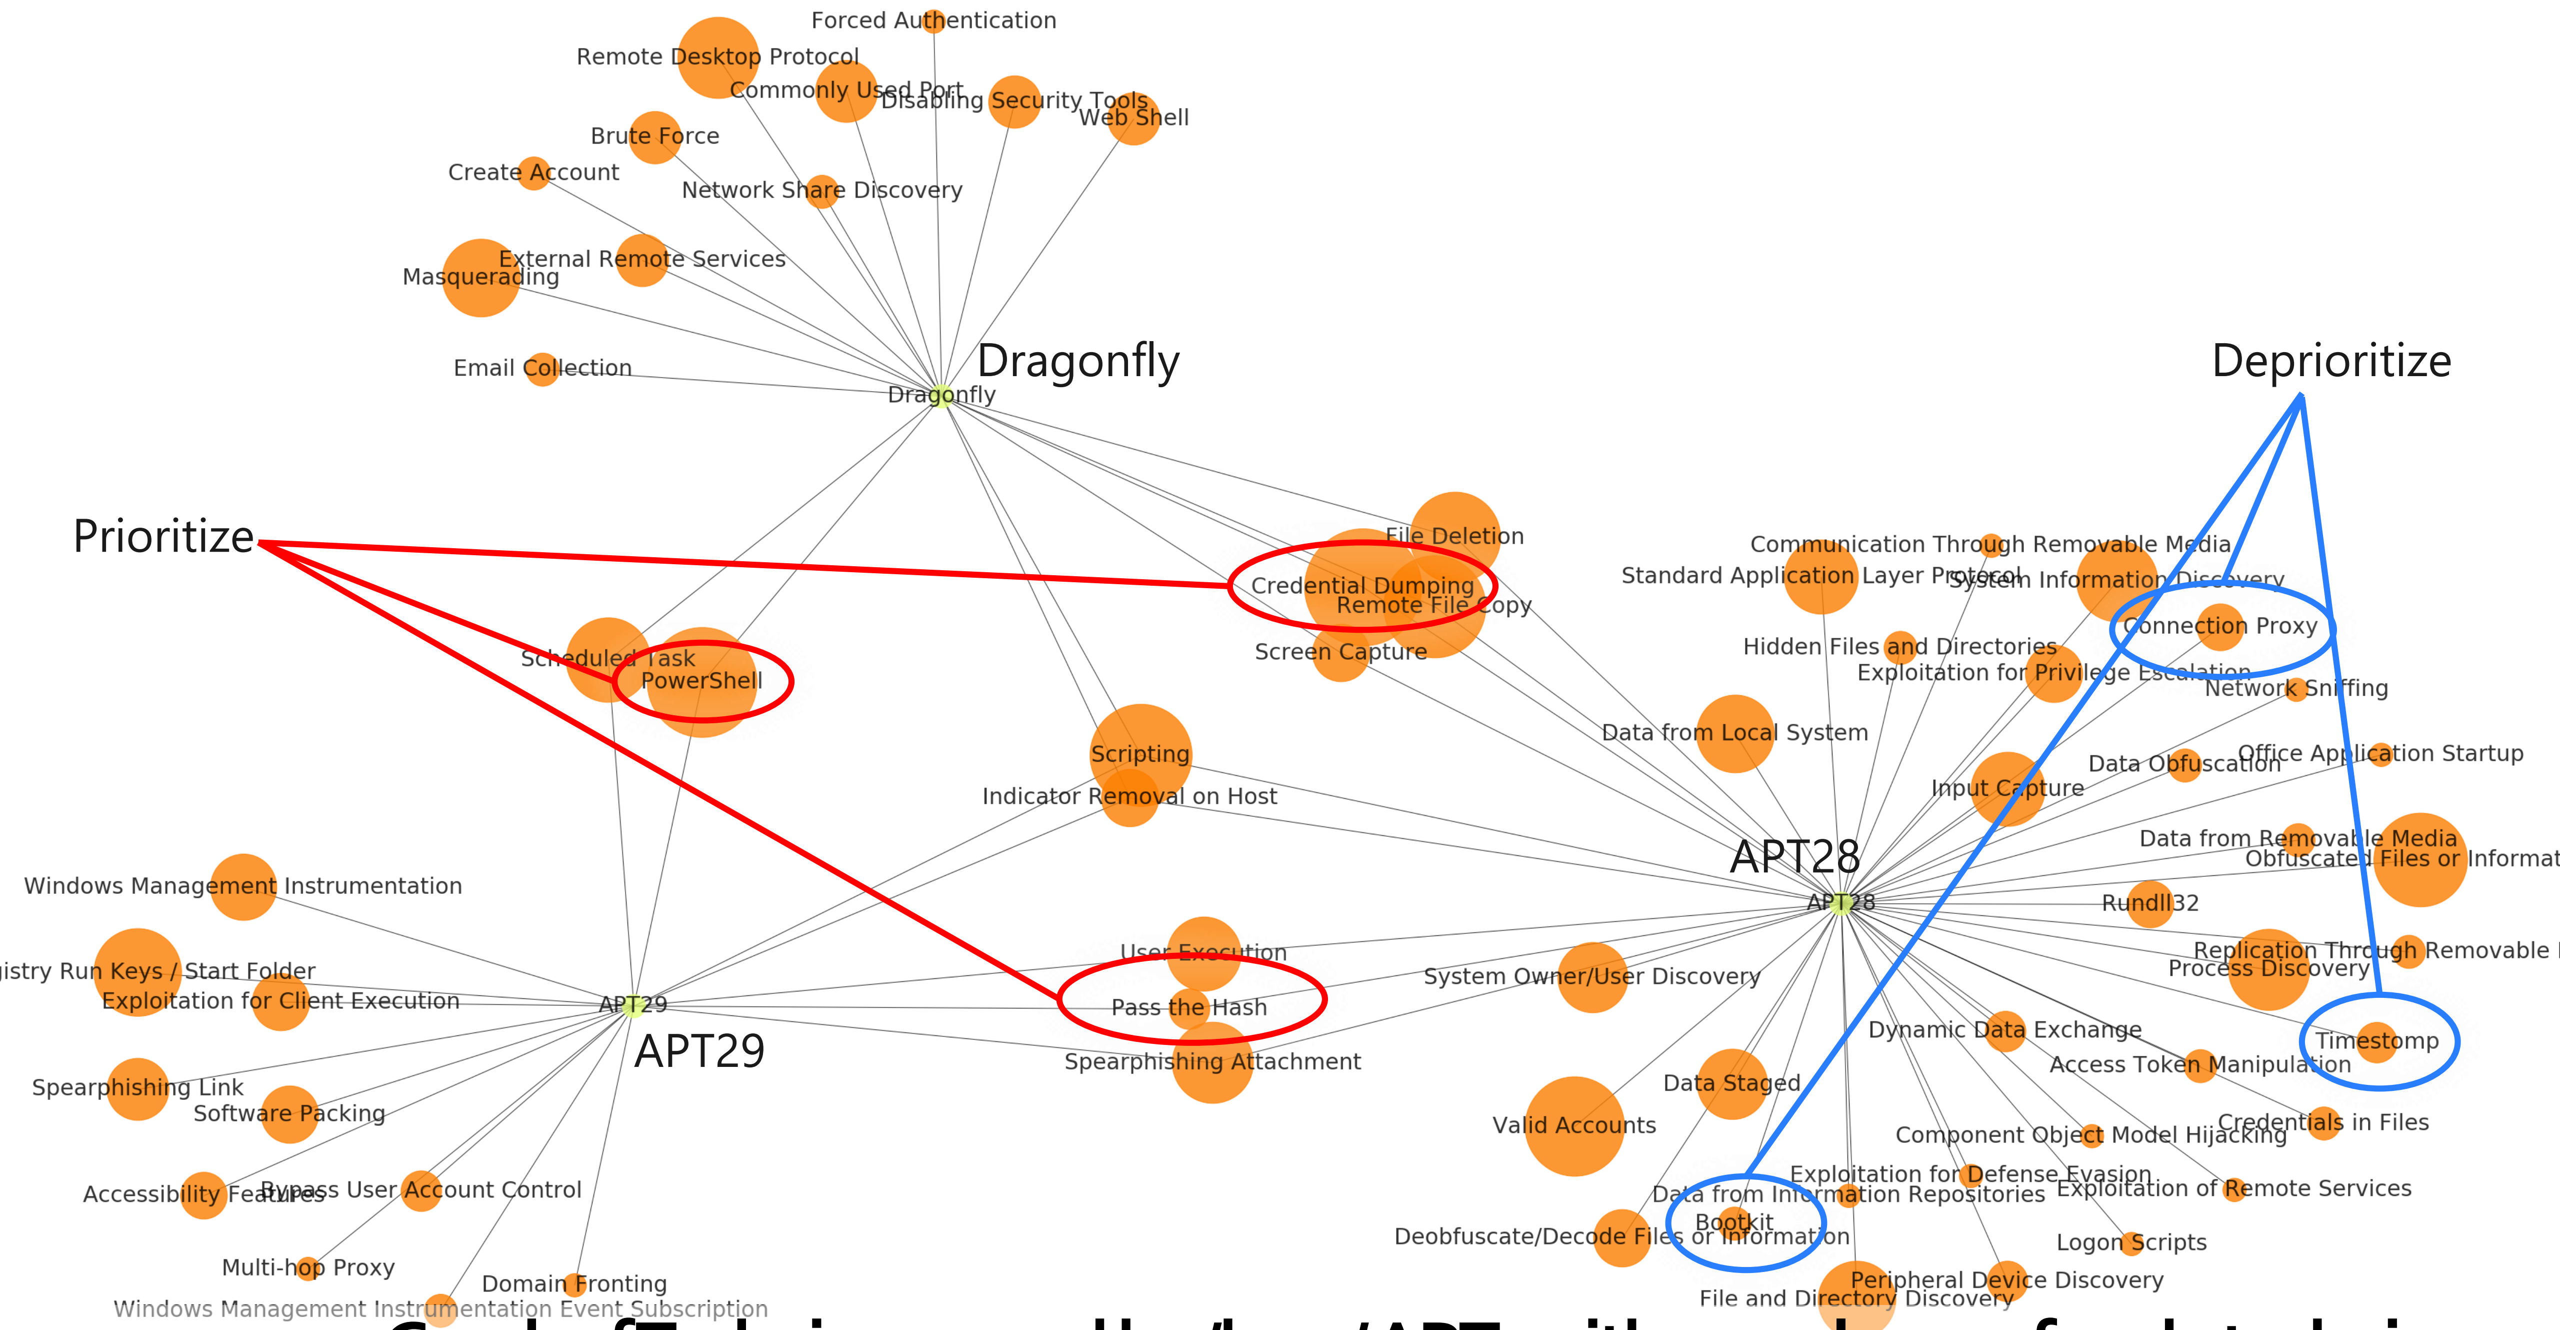

Graph of Techniques used by /bear/ APTs with prevalence of each technique

https://mitre-attack.github.io/caret/#/

Can this be done by machine learning?

Input

- Written material
- Blogs
- Whitepapers
- Incident Response reports

{black box}

- Extract actor names
- Extract tools names
- Extract techniques
- Extract relationships

Output

- Attacker graphs
- Timelines

Agenda

- Introduce the idea of Named Entity Extraction
- Build a machine learning/deep learning based Cyber Entity Extractor
  - Training Data
  - Feature Extraction
  - Architecture and Models
  - Evaluation
- Demo
- Driving Impact

What is Named Entity Extraction?

Sansa Stark

Fictional character

Sansa Stark is the eldest daughter of Eddard Stark of Winterfell and his wife Catelyn. She initially starts off with a very naive view of the world, but as time goes on and she and her family suffer one cruelty and betrayal after another, she becomes a more hardened and learned individual. Wikia

Sansa Stark: PERSON

Eddard Stark: PERSON

Catelyn: PERSON

Winterfell: ORGANIZATION (GEOPOLITICAL ENTITY)

https://gameofthrones.fandom.com/wiki/Sansa_Stark

APT REPORTS

The Dropping Elephant – aggressive cyber-espionage in the Asian region

GReAT Global Research & Analysis Team

By GReAT on July 8, 2016. 5:57 am

Dropping Elephant (also known as "Chinastrats" and "Patchwork") is a relatively new threat actor that is targeting a variety of high profile diplomatic and economic targets using a custom set of attack tools. Its victims are all involved with China's foreign relations in some way, and are generally caught through spear-phishing or watering hole attacks.

Dropping Elephant: BADACTOR

Chinastrats: BADACTOR

Patchwork: BADACTOR

spear-phishing: TECHNIQUE

watering hole: TECHNIQUE

https://securelist.com/the-dropping-elephant-actor/75328/

Training our own Cyber Entity Extractor

Training Data

Feature Extraction

Architecture

Assessments

Training Data

- APT Notes
- Public threat intelligence blogs collected since June 2018
- 2704 Documents
- On average, ~1% of the tokens are “interesting”

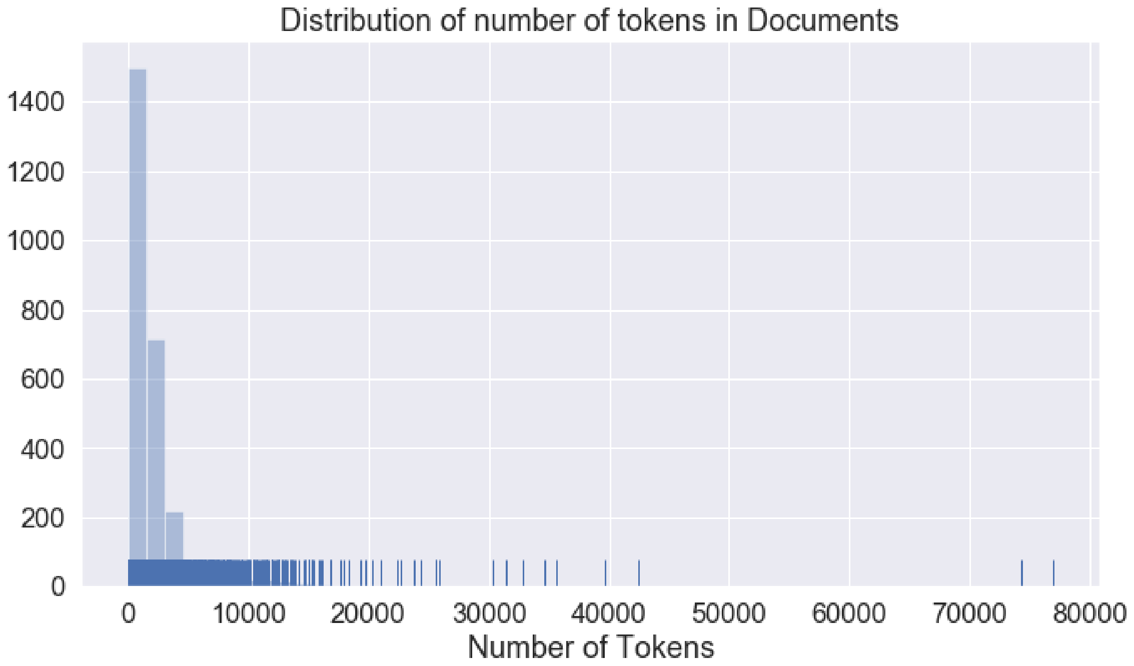

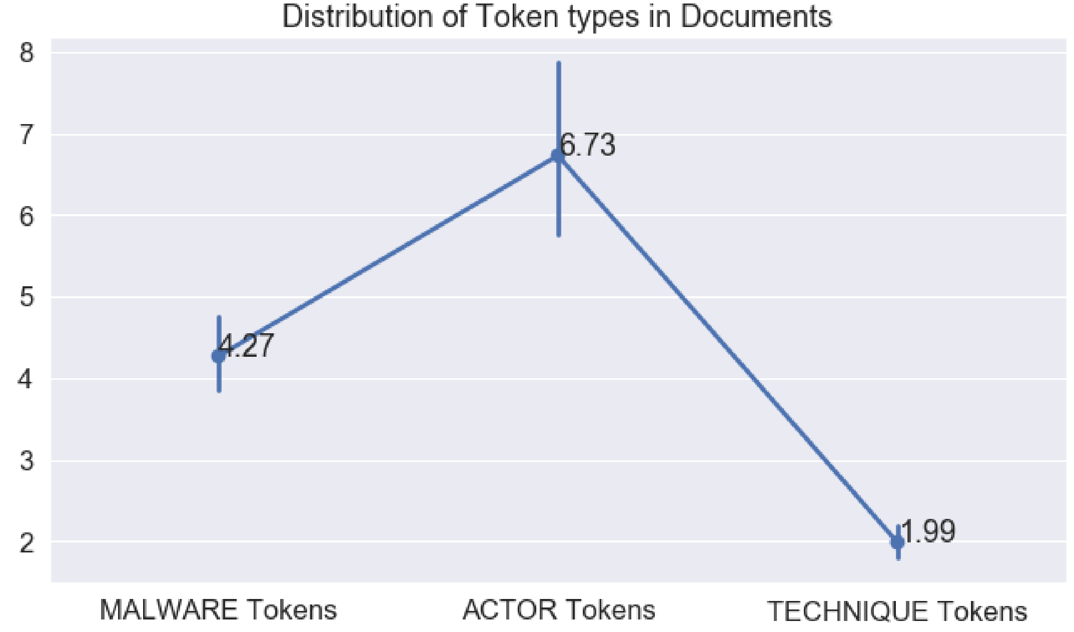

Training Data

- Labels: Caret Dataset (MITRE)

- Automatically annotated using longest extent pattern matching

- Kinda noisy, but best we can do short of manual annotation


```
  "name": "Group/G0007",
  "techniques": [
  "ID": "G0007",
  "aliases": [
    "APT28",
    "Sednit",
    "Sofacy",
    "Pawn Storm",
    "Fancy Bear",
    "STRONTIUM",
    "Tsar Team",
    "Threat Group-4127",
    "TG-4127"
  ]
},
{
  "name": "Group/G0016",
  "techniques": [
  "ID": "G0016",
  "aliases": [
    "APT29",
    "The Dukes",
    "Cozy Bear",
    "CozyDuke"
  ]
},
{
  "name": "Group/G0022",
  "techniques": [
  "ID": "G0022",
  "aliases": [
    "APT3",
    "Gothic Panda",
    "Pirpi",
    "UPS Team",
    "Buckeye",
    "Threat Group-0110",
    "TG-0110"
  ]
```

Training Data

Numbered Panda (also known as IXESHE, DynCalc, DNSCALC, and APT12) is a cyber espionage group believed to be linked with the Chinese military.

```
nltk.tree2conlltags(nltk.ne_chunk(nltk.pos_tag(nltk.word_tokenize(sansa))))
```

```
[('Sansa', 'NNP', 'B-PERSON'),
 ('Stark', 'NNP', 'I-PERSON'),
 ('is', 'VBZ', 'O'),
 ('the', 'DT', 'O'),
 ('eldest', 'JJS', 'O'),
 ('daughter', 'NN', 'O'),
 ('of', 'IN', 'O'),
 ('Eddard', 'NNP', 'B-PERSON'),
 ('Stark', 'NNP', 'I-PERSON'),
 ('of', 'IN', 'O'),
 ('Winterfell', 'NNP
 ('and', 'CC', 'O'),
 ('his', 'PRP$', 'O'
 ('wife', 'NN', 'O'),
 ('Catelyn', 'NNP', 'B-PERSON'),
 ('.', '.', 'O'),
 ('She', 'PRP', 'O'),
 ('initially', 'RB', 'O'),
 ('starts', 'VBZ', 'O'),
```

IOB Style

```
('Eddard', 'NNP', 'B-PERSON'),
('Stark', 'NNP', 'I-PERSON'),
```

('Numbered', 'B-BADACTOR'),
('Panda', 'I-BADACTOR'),
('(', 'O'),
('also', 'O'),
('known', 'O'),
('as', 'O'),
('IXESHE', 'B-BADACTOR'),
(',', 'O'),
('DynCalc', 'B-BADACTOR'),
(',', 'O'),
('DNSCALC', 'B-BADACTOR'),
(',', 'O'),
('and', 'O'),
('APT12', 'B-BADACTOR'),
(')', 'O'),
('is', 'O'),
('a', 'O'),
('cyber', 'O'),
('espionage', 'O'),
('group', 'O'),
('believed', 'O'),
('to', 'O'),
('be', 'O'),
('linked', 'O'),
('with', 'O'),
('the', 'O'),
('Chinese', 'O'),
('military', 'O'),
('.', 'O')

Training our own Cyber Entity Extractor

Training Data

Feature Extraction

Architecture

Assessments

Feature Extraction

- Traditional
  - Word itself
  - Part of speech
  - Lemma
  - Word type (alphanumeric, digits, punctuation)
  - Orthographic features (lowercase, ALLCAPS, Upper initial, MiXedCaPs etc.)
- Unsupervised
  - Word Embeddings

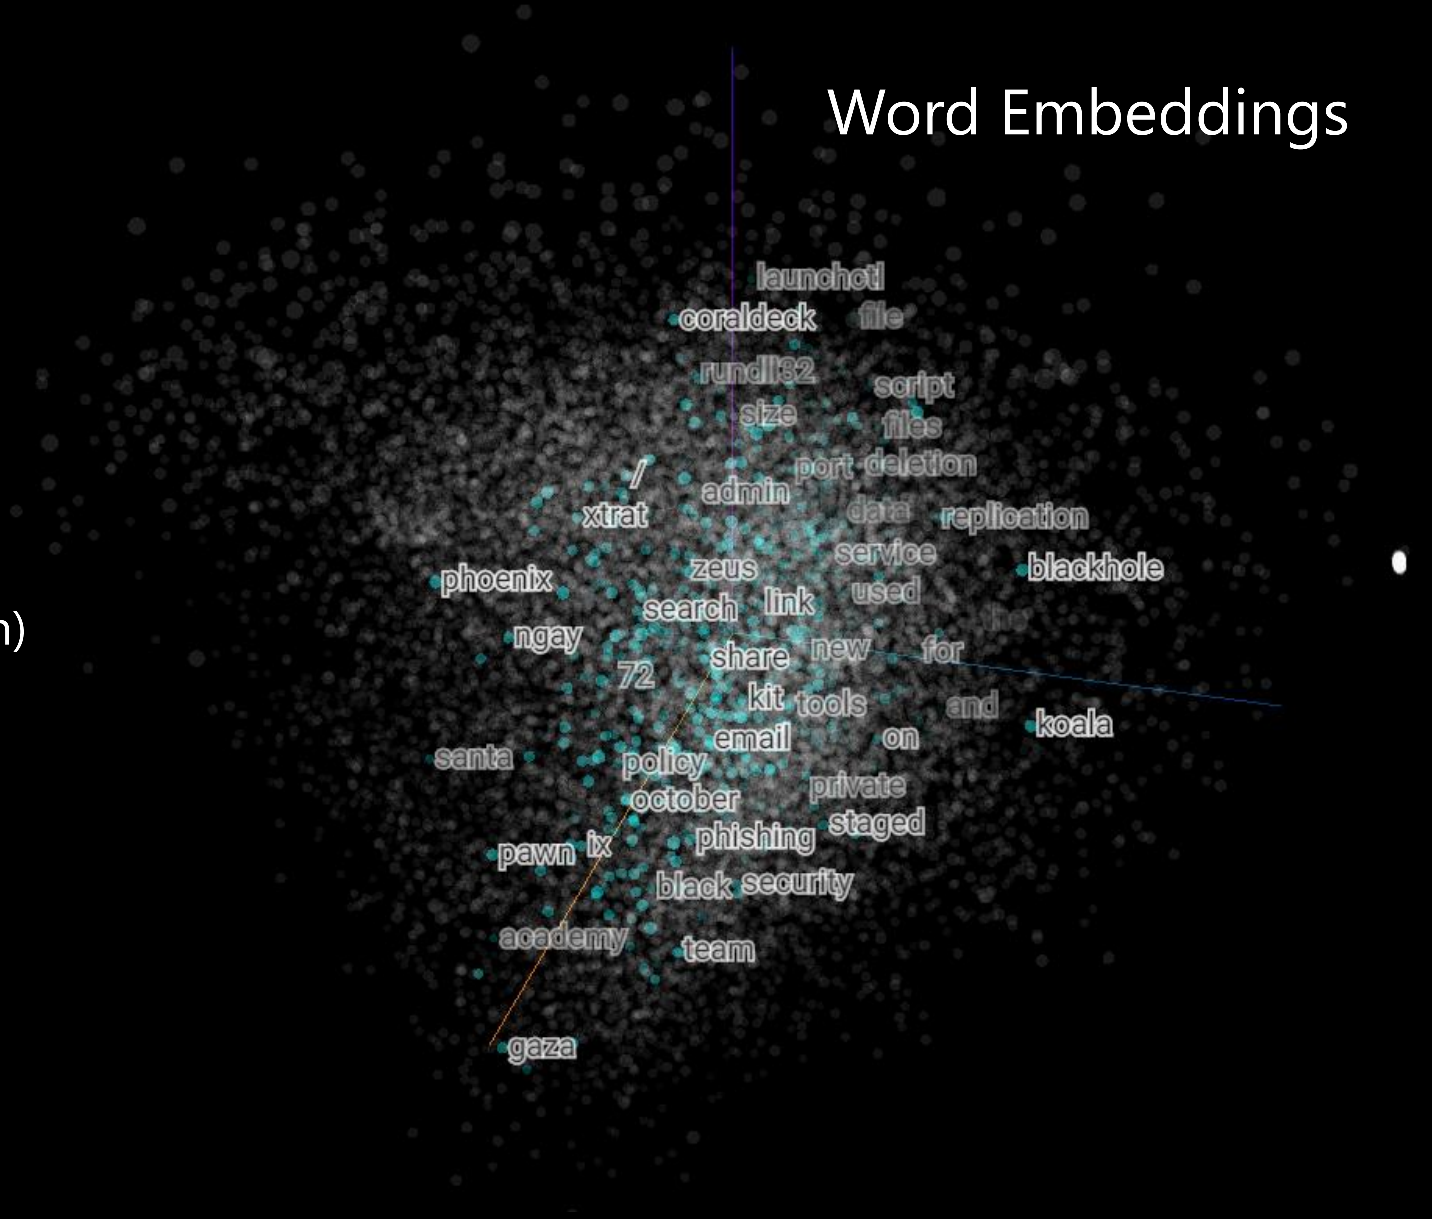

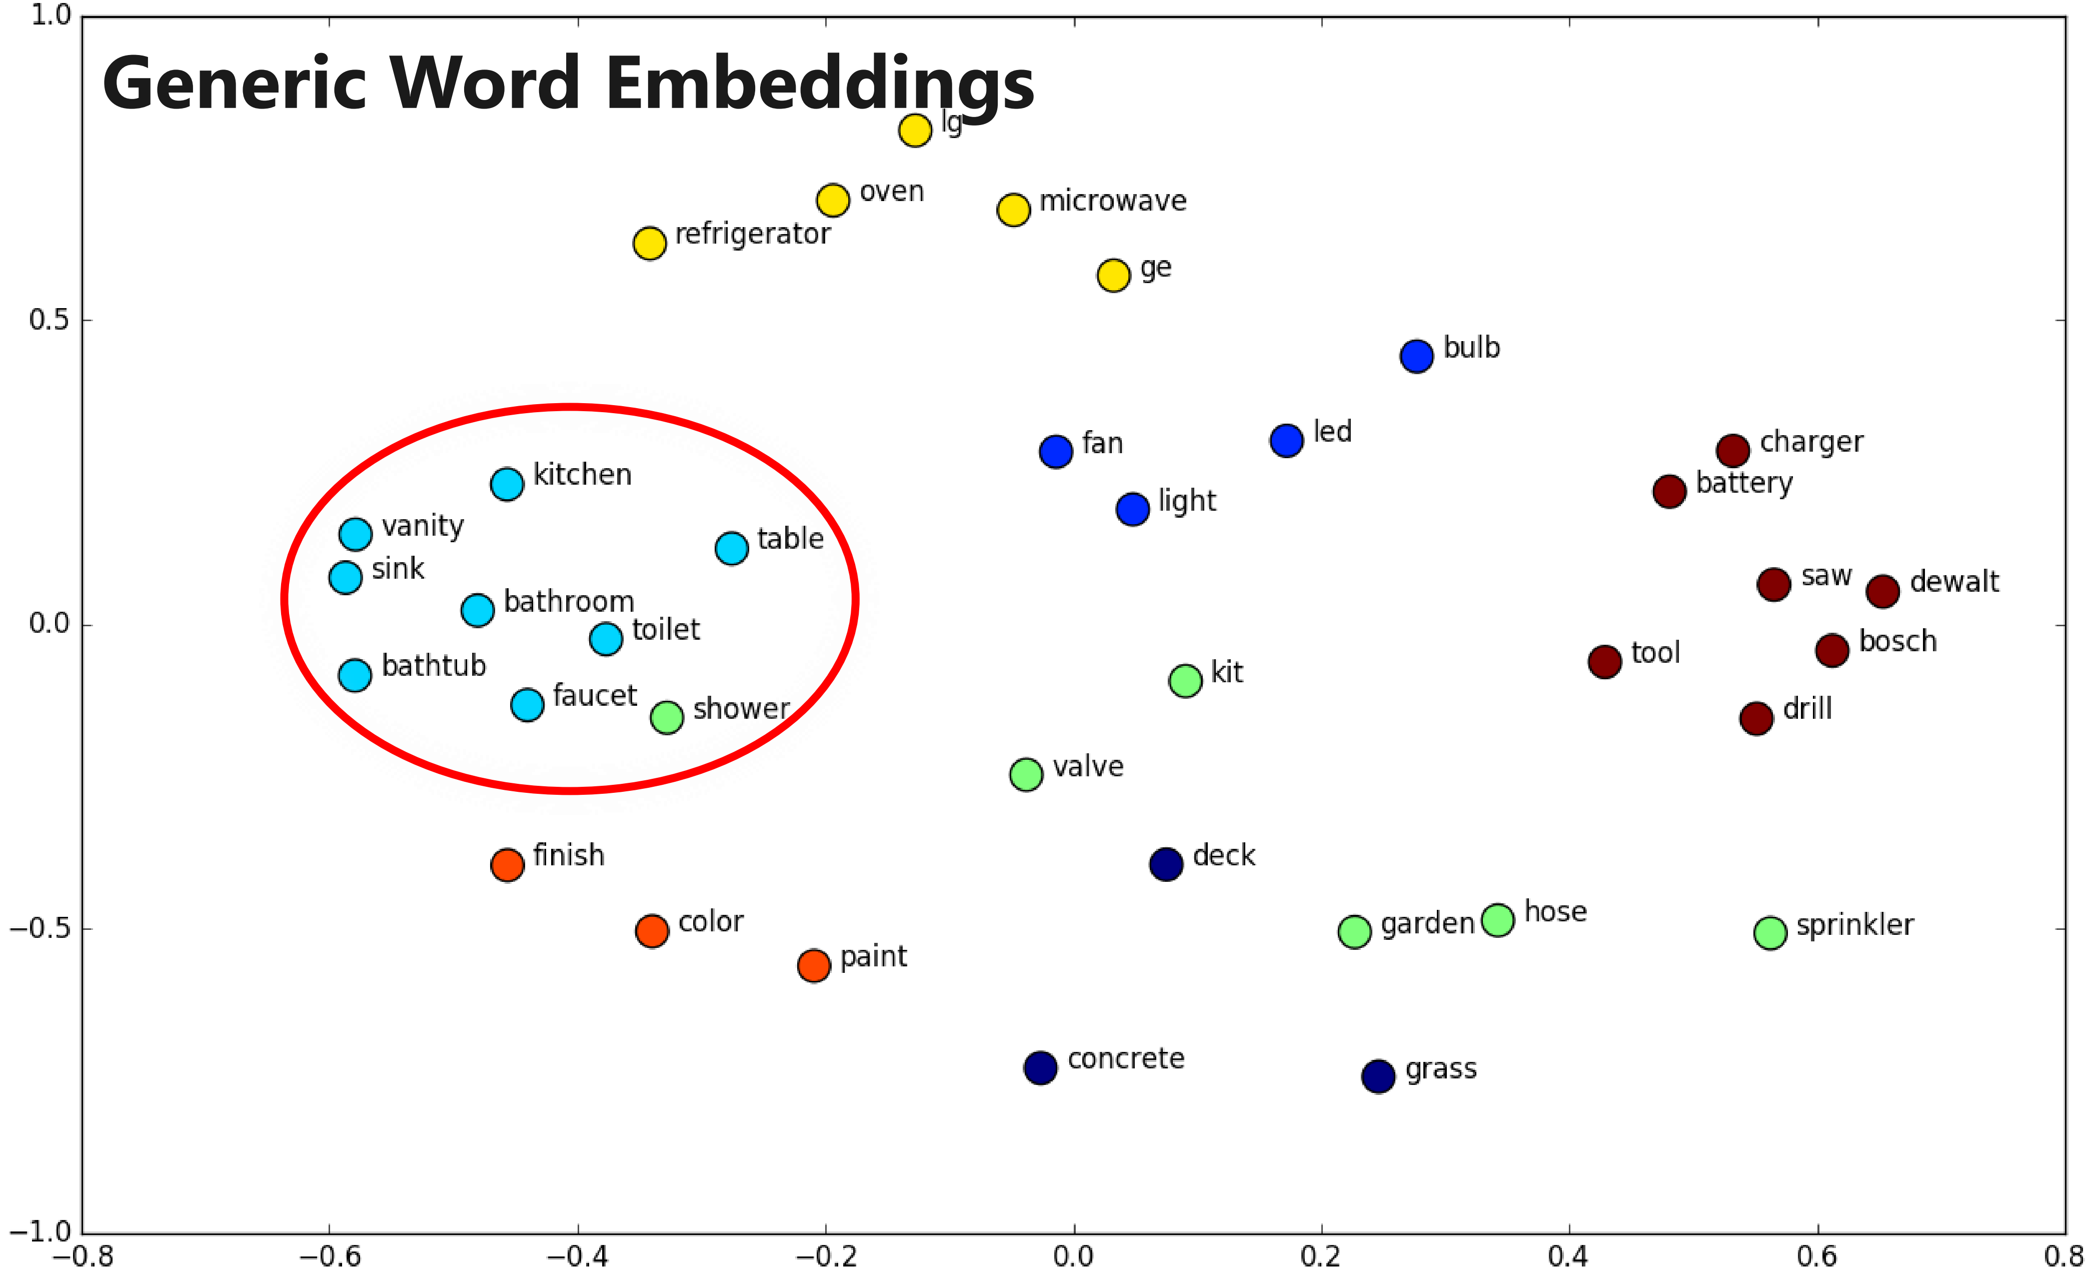
Generic Word Embeddings
lg
oven
microwave
refrigerator
ge
bulb
fan
led
charger
battery
kitchen
light
vanity
table
sink
saw
dewalt
bathroom
toilet
tool
bosch
bathtub
kit
faucet
shower
drill
valve
finish
deck
color
hose
garden
sprinkler
paint
concrete
grass
1.0
0.5
0.0
−0.5
−1.0
−0.8
−0.6
−0.4
−0.2
0.0
0.2
0.4
0.6
0.8

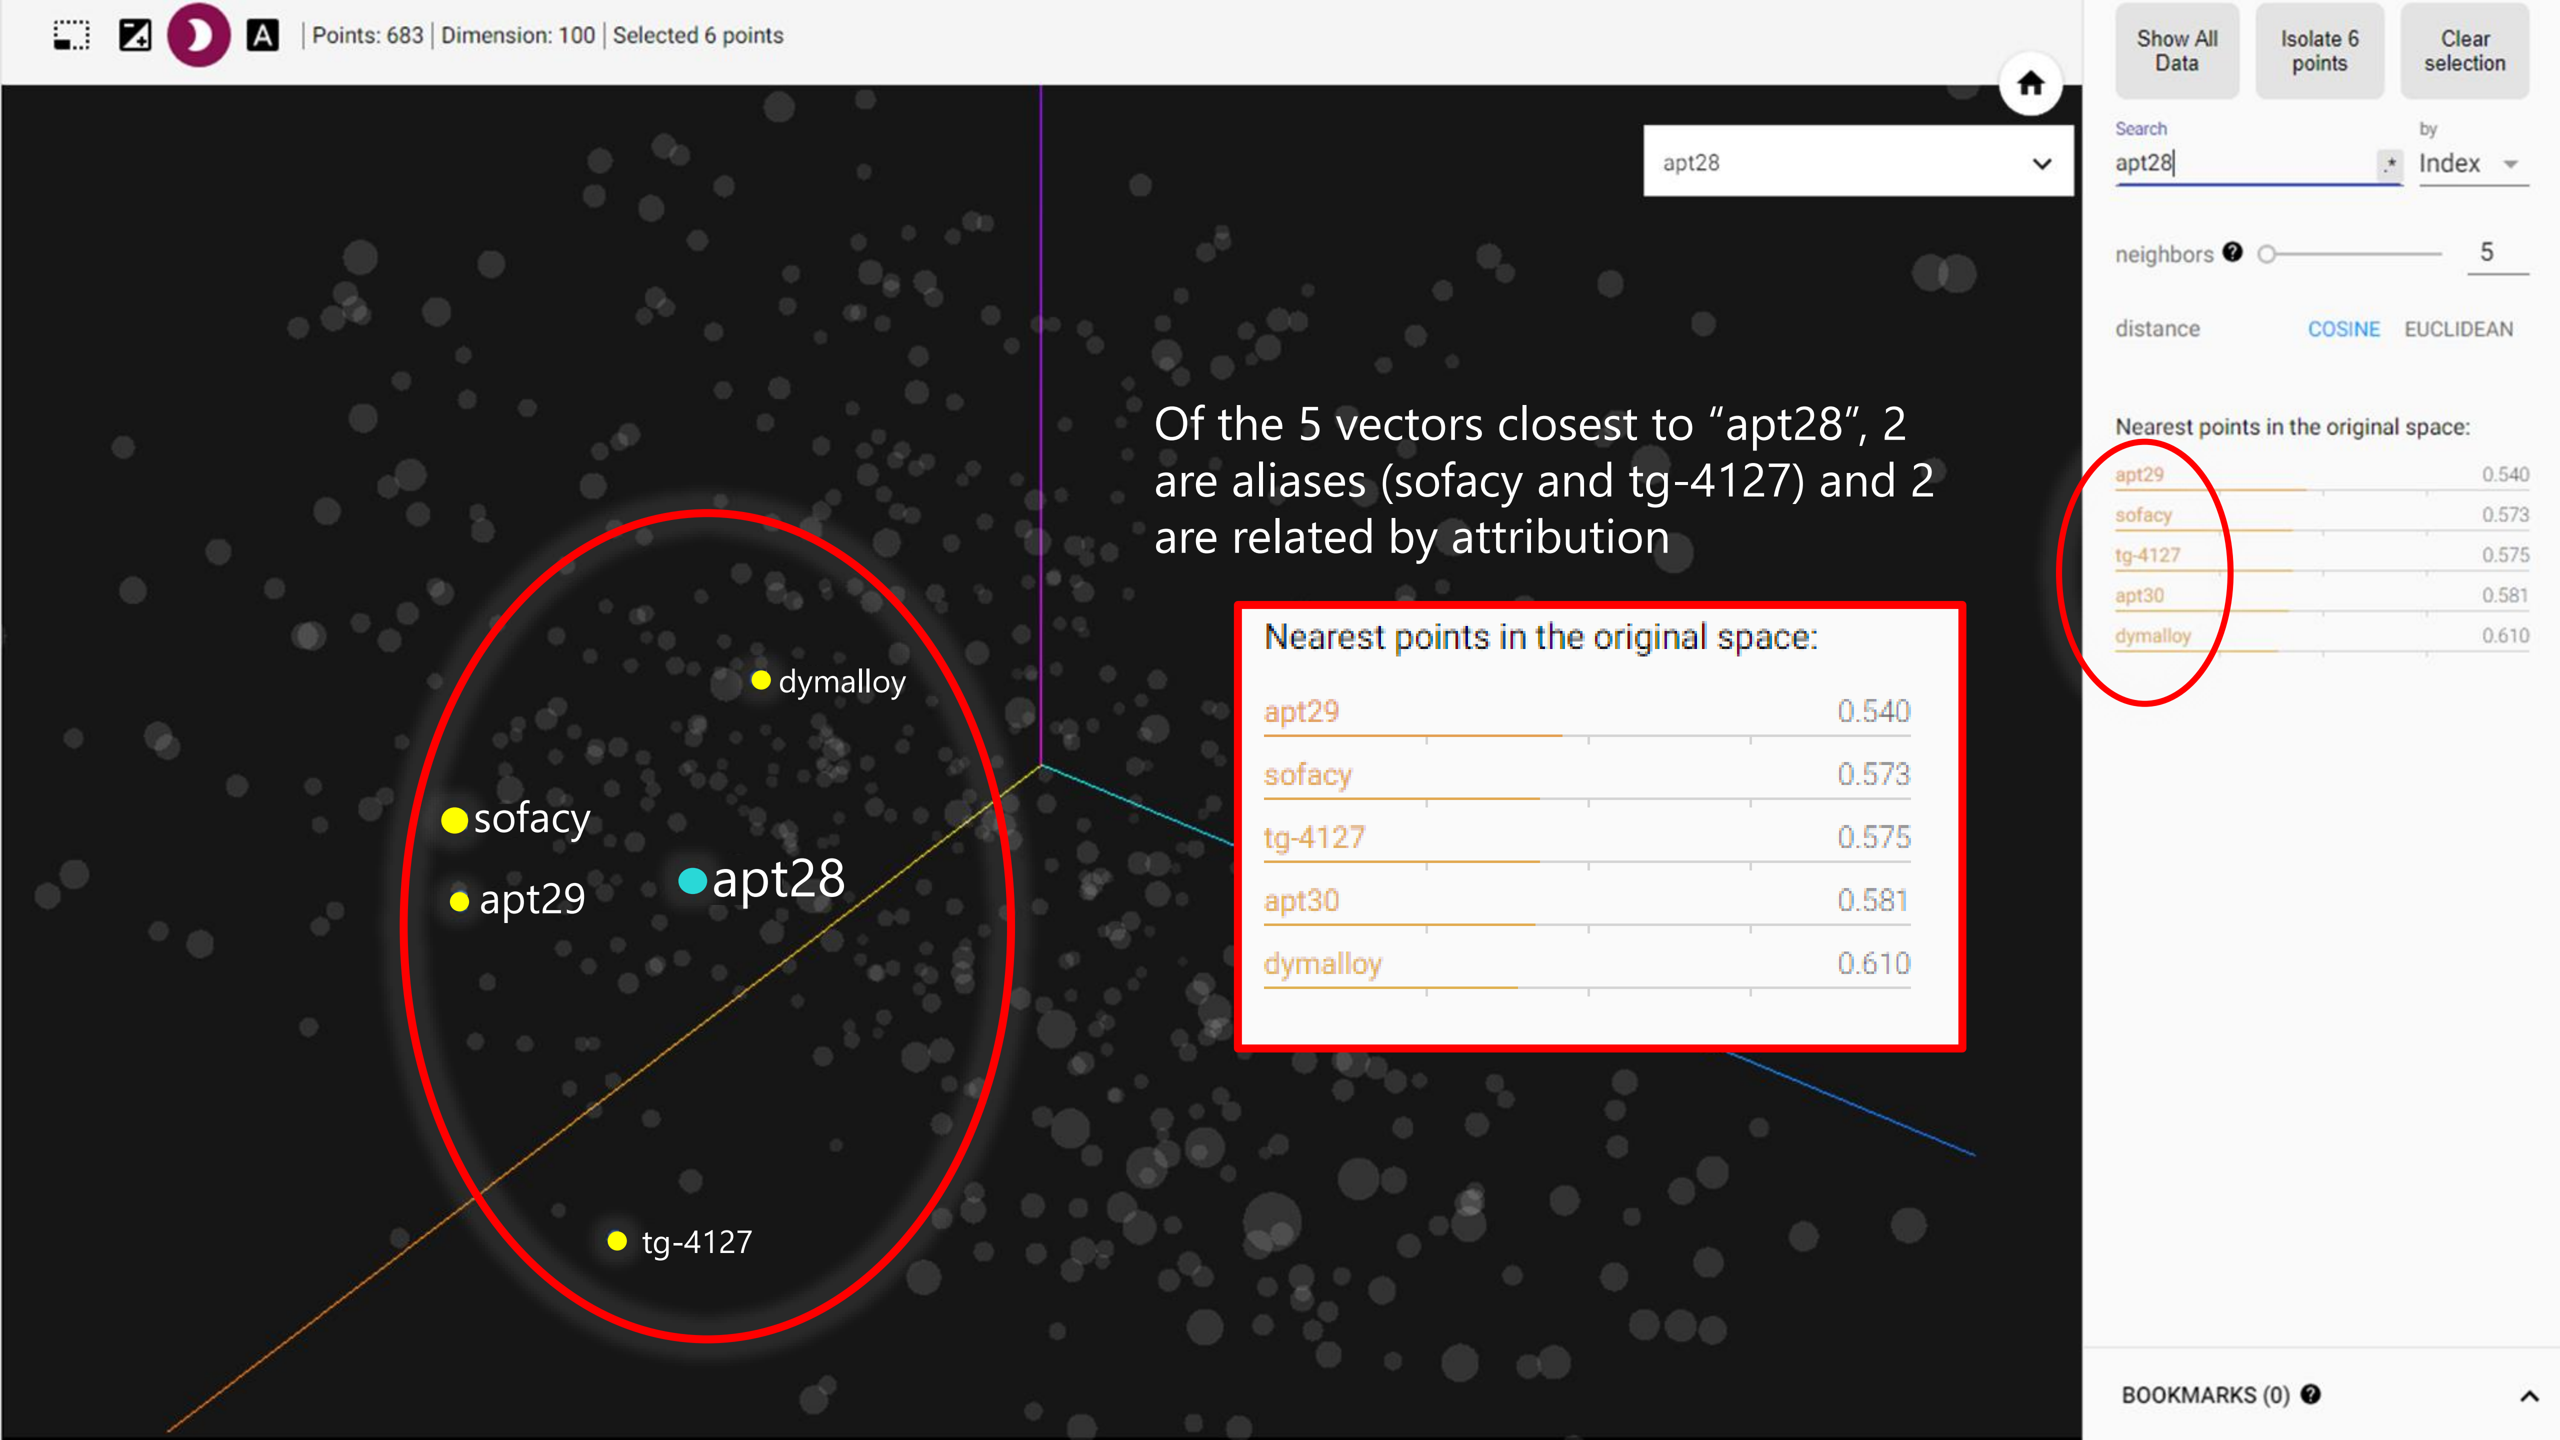

Points: 683 | Dimension: 100 | Selected 6 points
Show All Data
Isolate 6 points
Clear selection
apt28
Search
apt28
by
Index
neighbors
5
distance
COSINE
EUCLIDEAN
Nearest points in the original space:
apt29 0.540
sofacy 0.573
tg-4127 0.575
apt30 0.581
dymalloy 0.610
Of the 5 vectors closest to "apt28", 2 are aliases (sofacy and tg-4127) and 2 are related by attribution
Nearest points in the original space:
apt29 0.540
sofacy 0.573
tg-4127 0.575
apt30 0.581
dymalloy 0.610
dymalloy
sofacy
apt29
apt28
tg-4127
BOOKMARKS (0)

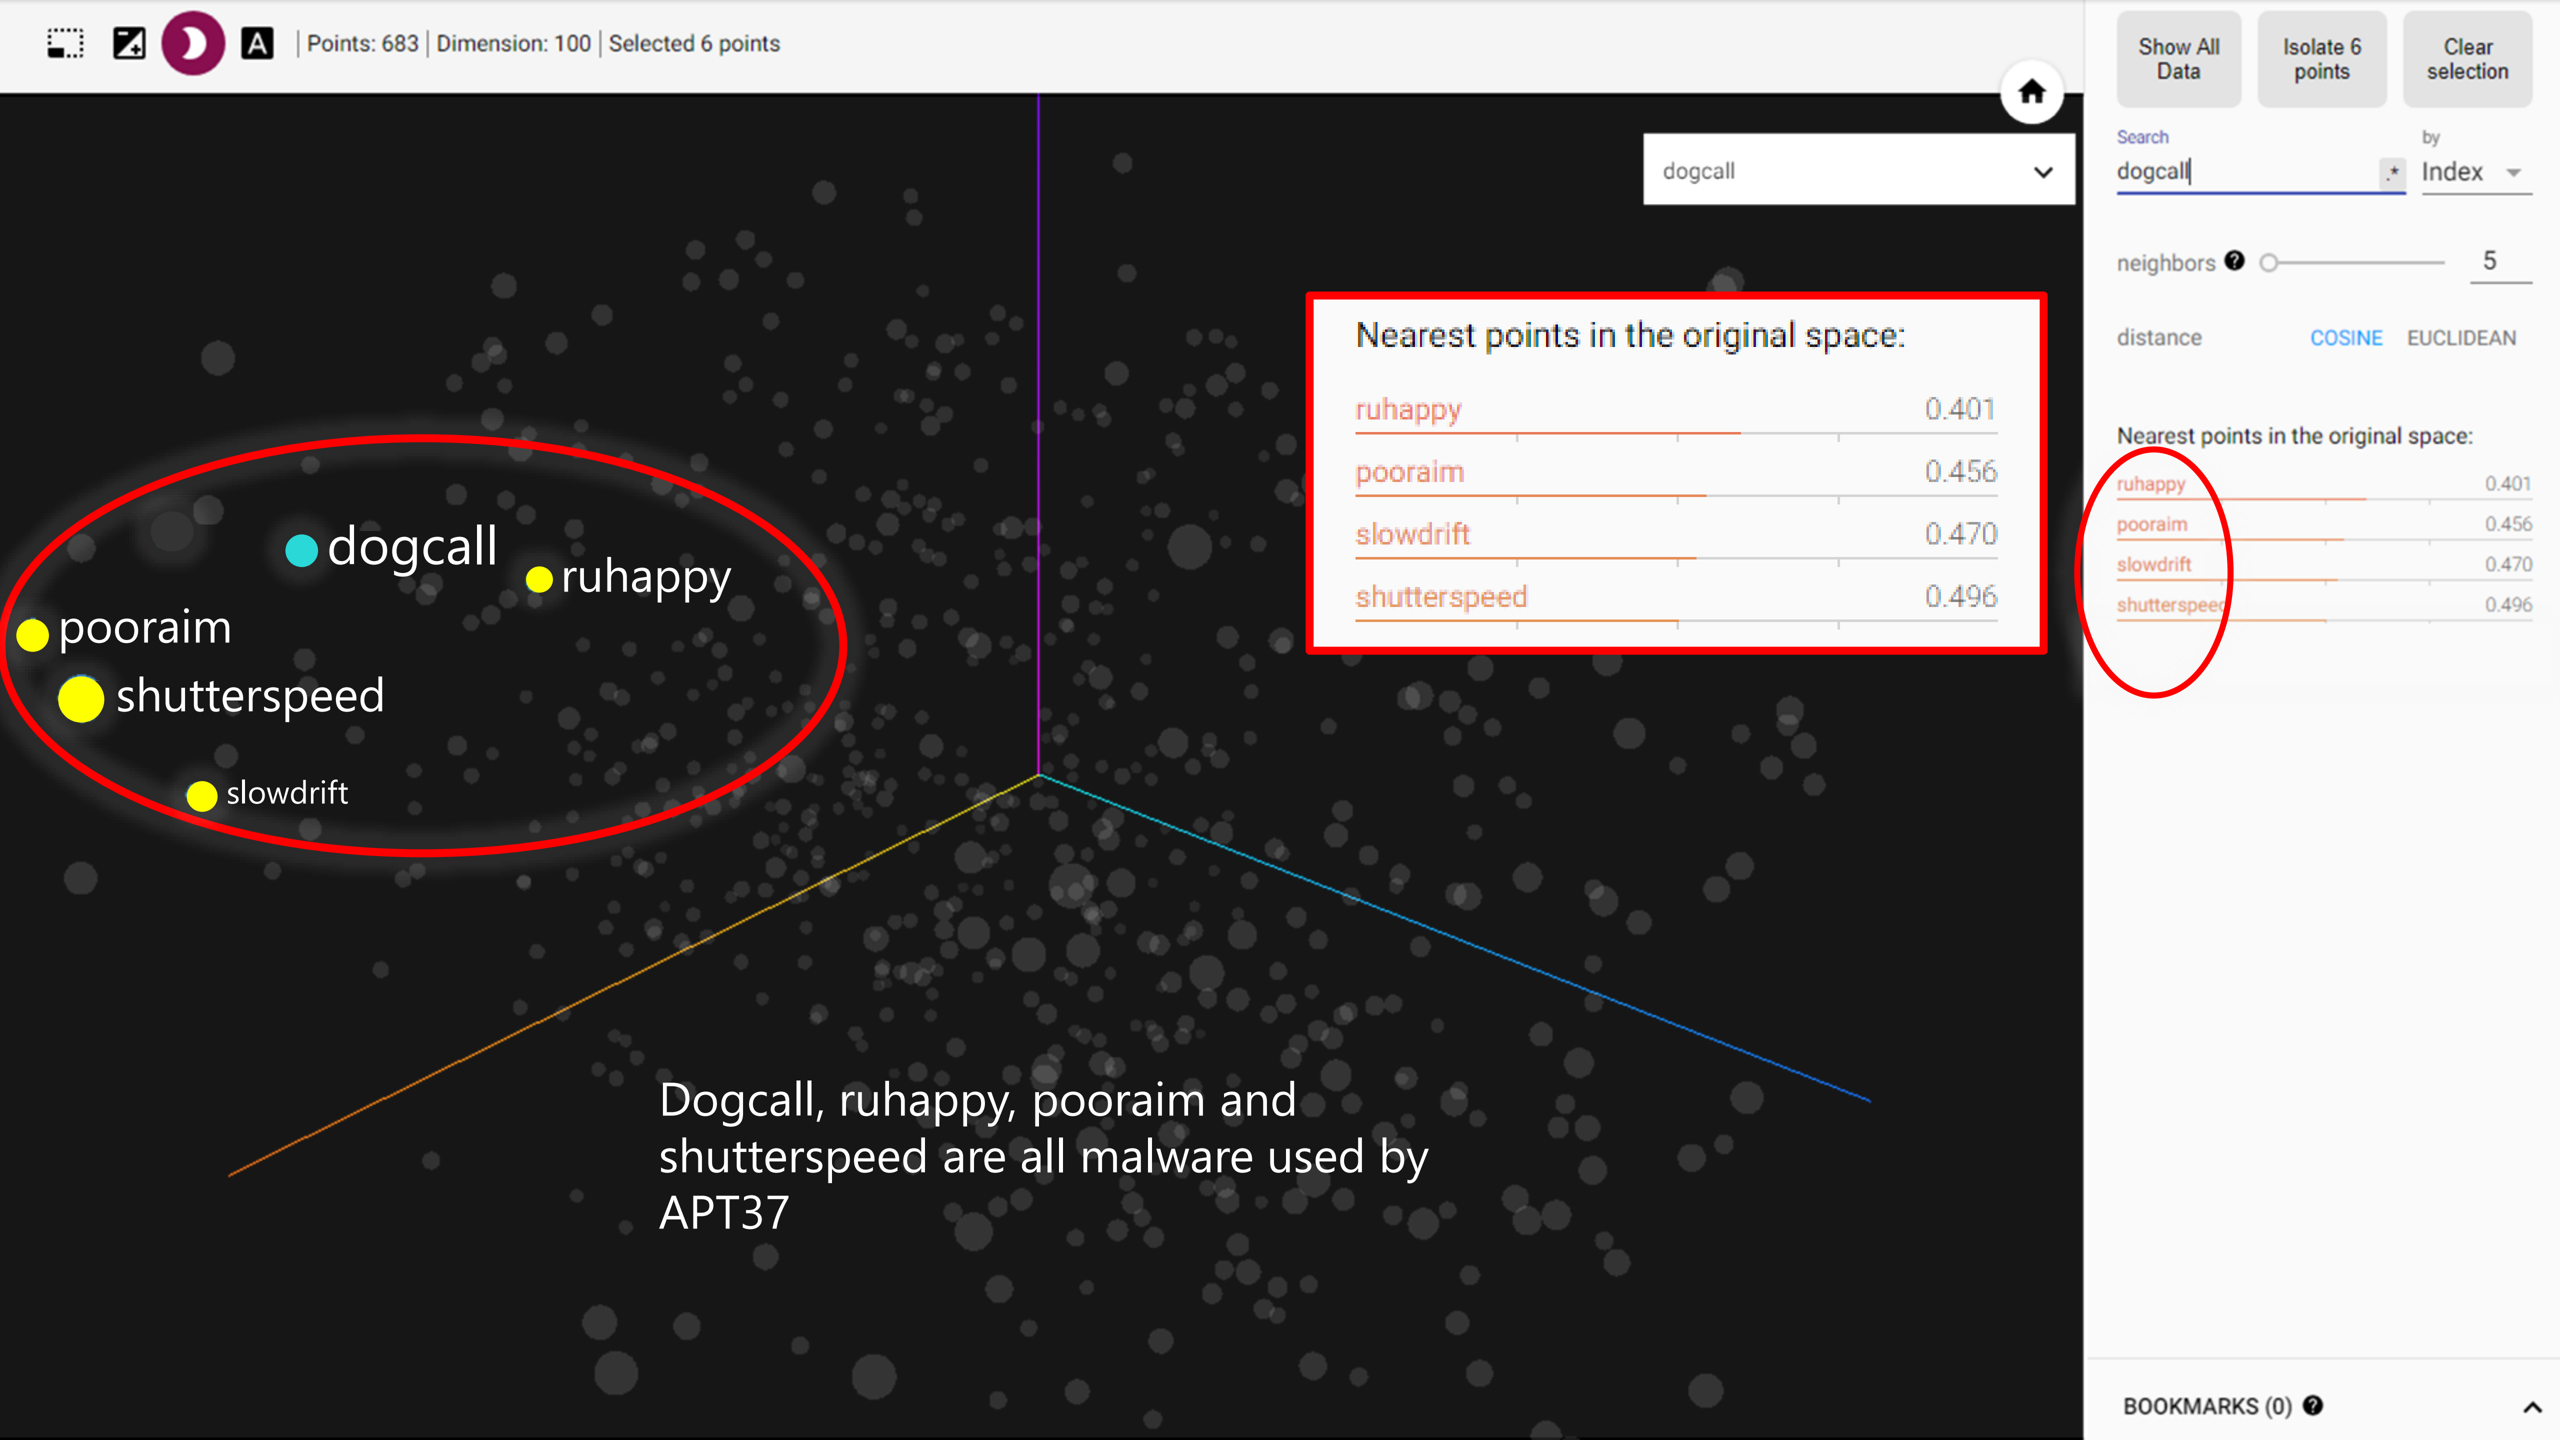
Points: 683 | Dimension: 100 | Selected 6 points
Show All Data
Isolate 6 points
Clear selection
Search
dogcall
by
Index
neighbors
5
distance
COSINE
EUCLIDEAN
Nearest points in the original space:
ruhappy 0.401
pooraim 0.456
slowdrift 0.470
shutterspeed 0.496
BOOKMARKS (0)
dogcall
dogcall
ruhappy
pooraim
shutterspeed
slowdrift
Nearest points in the original space:
ruhappy 0.401
pooraim 0.456
slowdrift 0.470
shutterspeed 0.496
Dogcall, ruhappy, pooraim and shutterspeed are all malware used by APT37

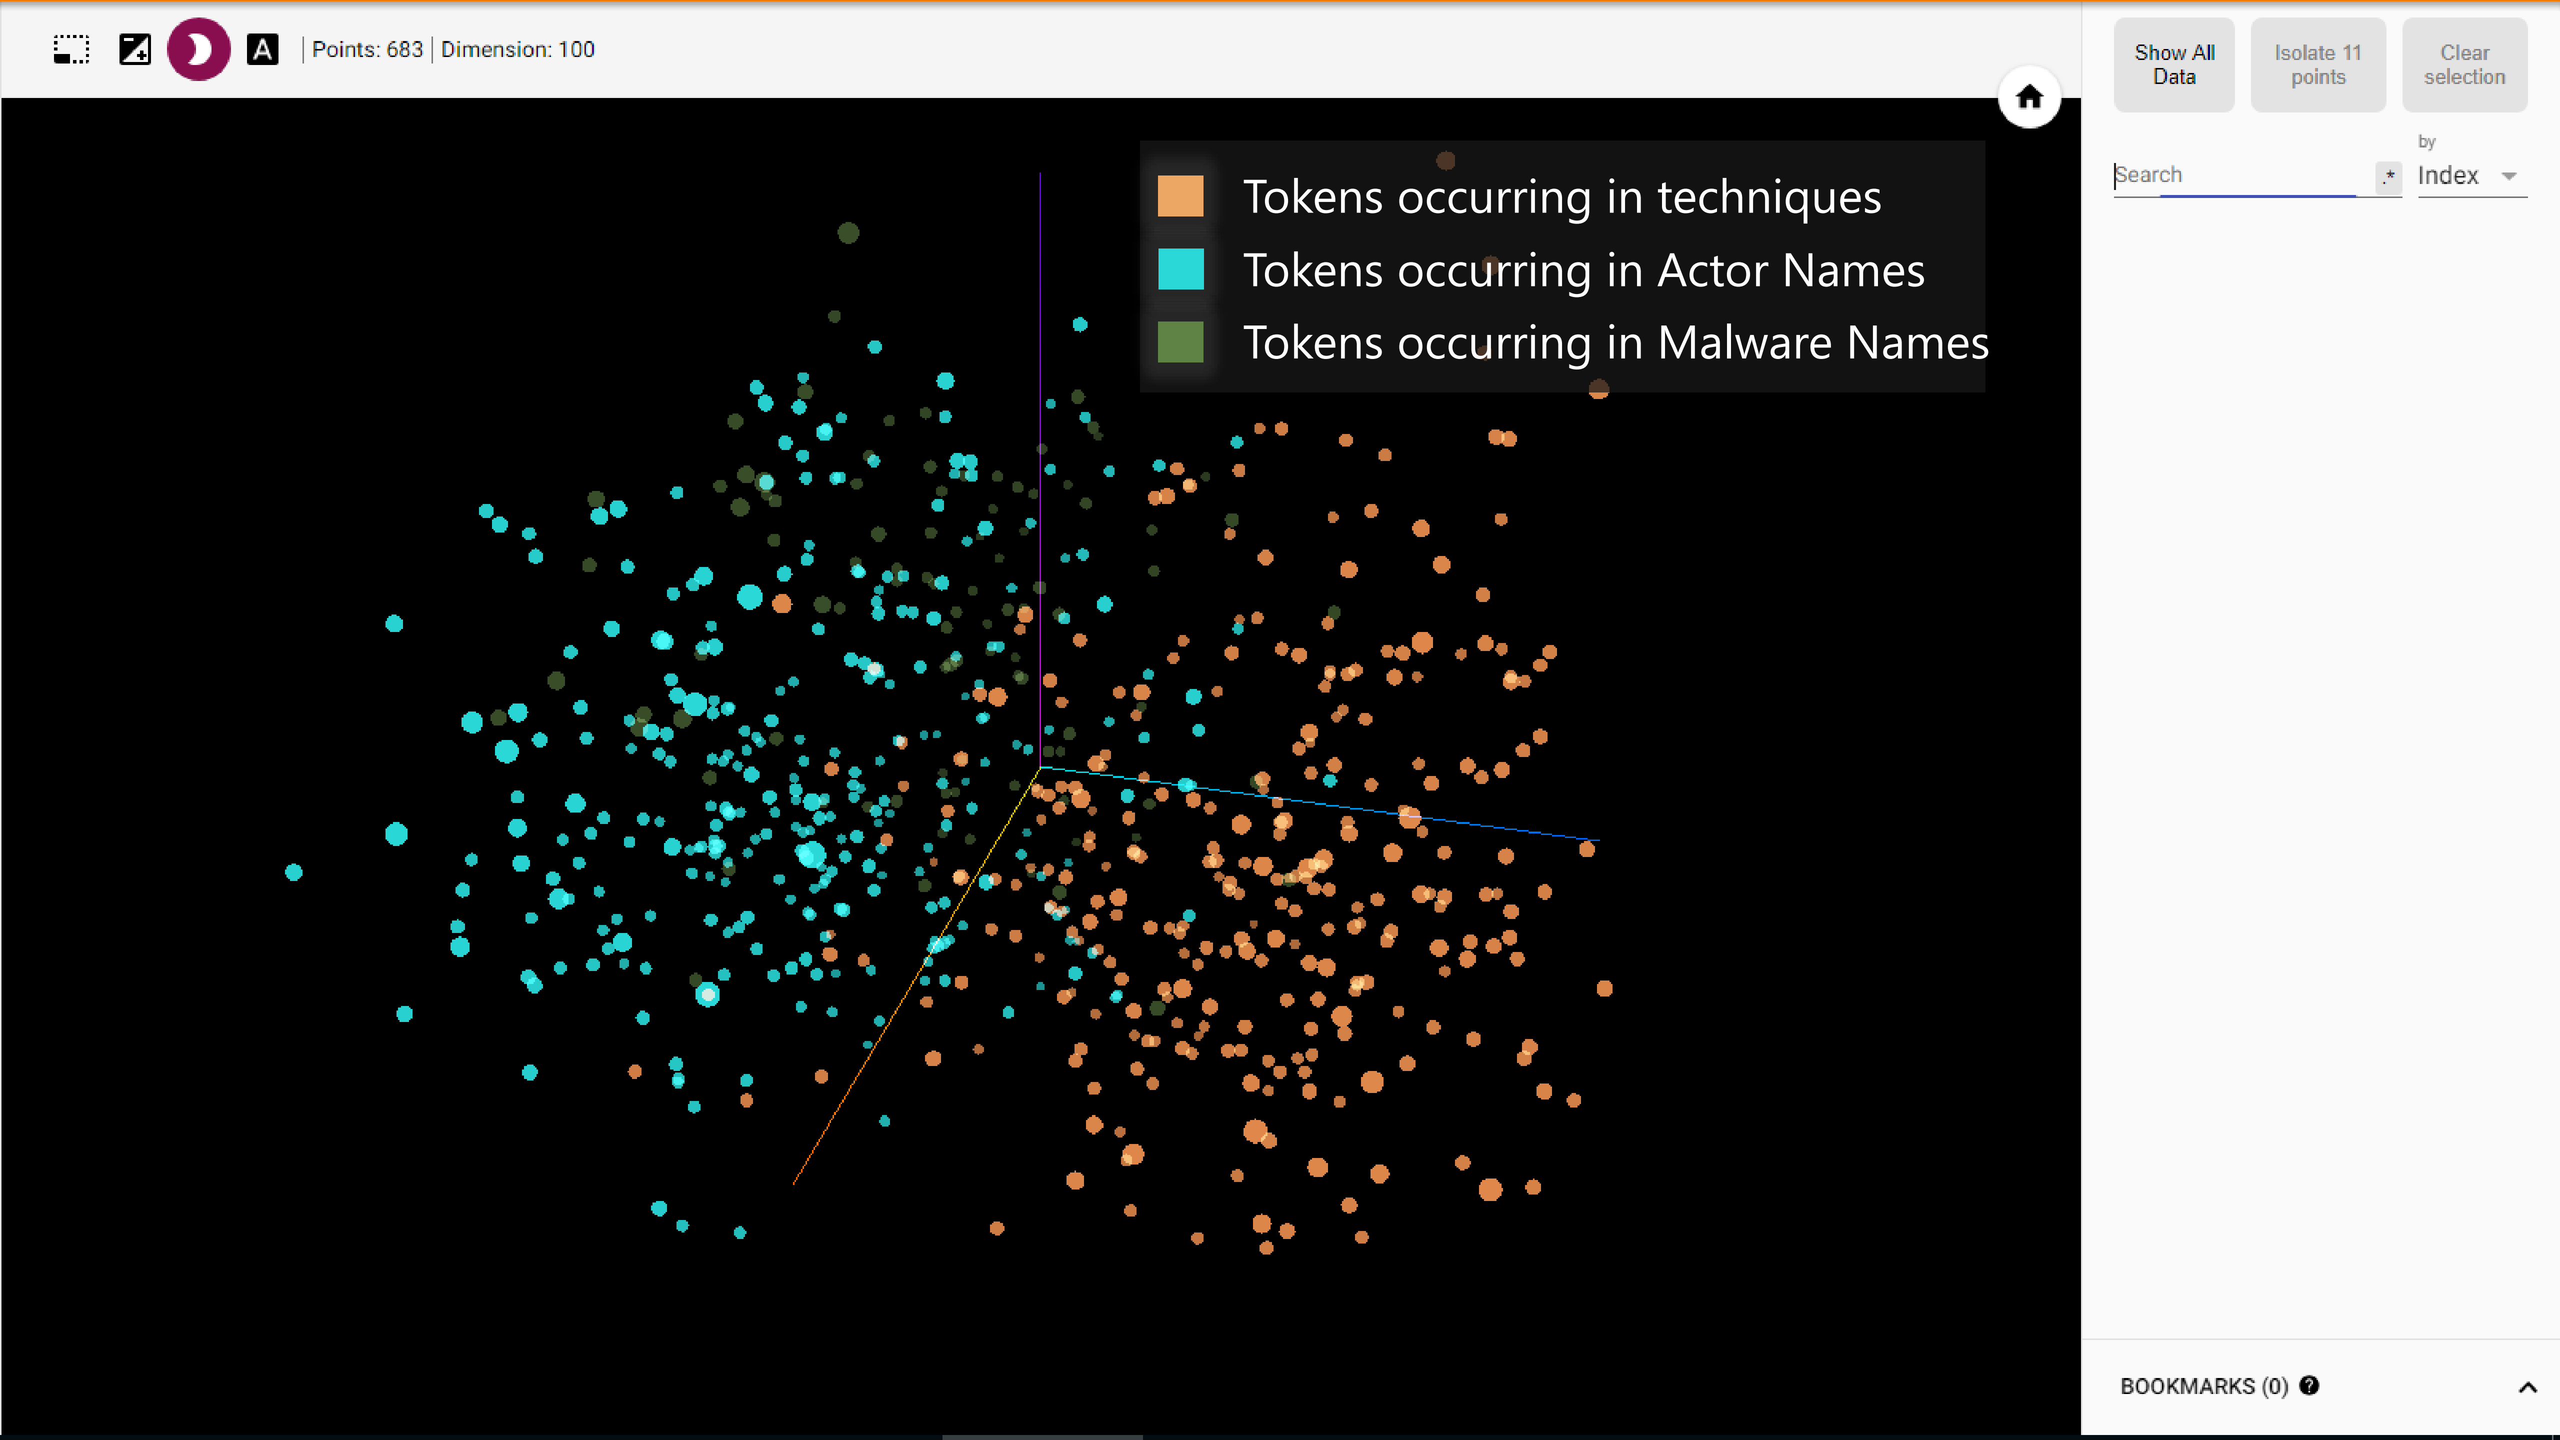

Points: 683 | Dimension: 100
Show All Data
Isolate 11 points
Clear selection
by
Search
Index
Tokens occurring in techniques
Tokens occurring in Actor Names
Tokens occurring in Malware Names
BOOKMARKS (0)

Training our own Cyber Entity Extractor

Training Data

Feature Extraction

Architecture

Assessments

Architecture

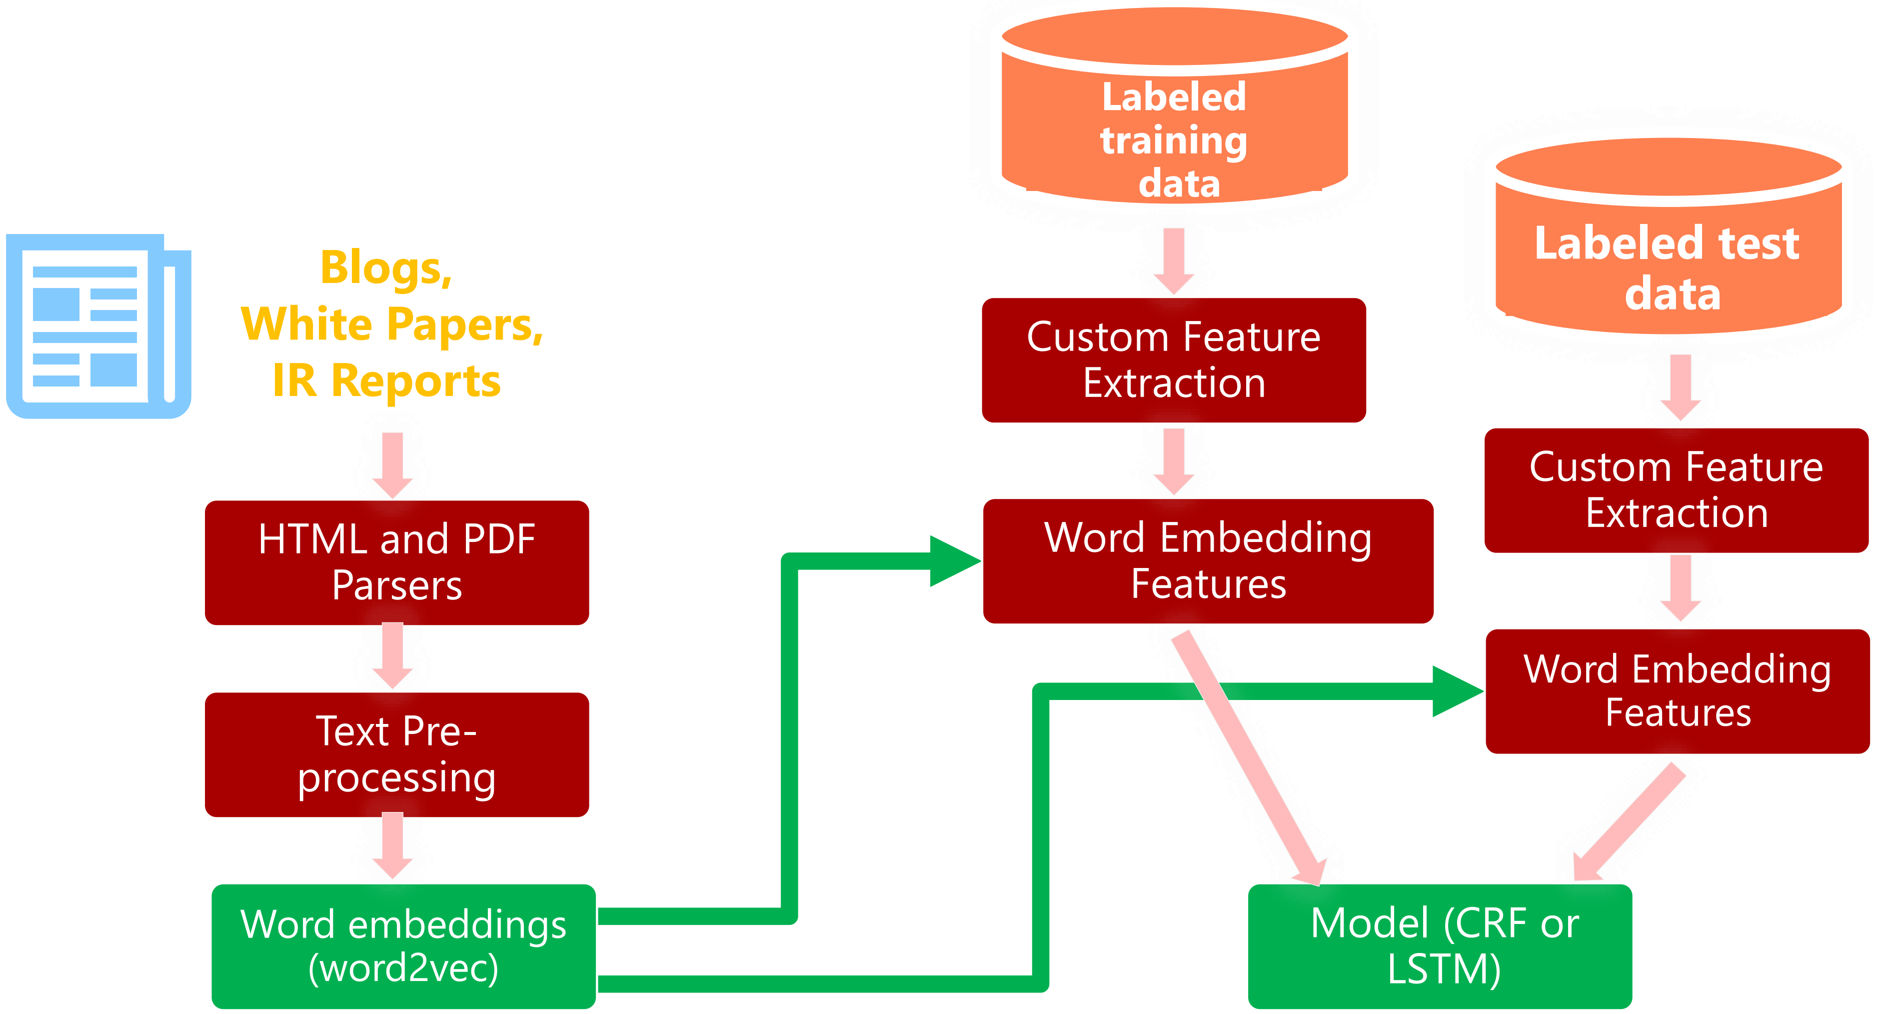

Conditional Random Fields (CRF)

- Statistical modeling method
- Not Deep Learning
- Used for sequence labeling tasks.

From \ To	O	B-BADACTOR	I-BADACTOR	B-MALWARE	I-MALWARE
O	2.742	0.22	-5.811	-0.084	-4.525
B-BADACTOR	-0.12	-1.071	2.568	-1.619	-0.642
I-BADACTOR	-0.176	-0.574	0.0	0.0	0.0
B-MALWARE	-0.253	-1.242	-1.391	-1.901	2.083
I-MALWARE	0.001	0.0	0.0	0.0	0.0

- Commonly used in Natural Language processing, biological sequences and computer vision
- Has short term memory
- 2 Experiments with CRF (one with and one without the word embeddings)

Long Short-term Memory (LSTM)

- Special type of RNN
- 2 Stacked Bidirectional LSTM Layers
- With Dropout
- Categorical Cross Entropy Loss Function
- Softmax activation for the final layer
- Keras + tensorflow

Embedding Layer	Output Shape:	None, 75, 100
Bidirectional (LSTM)	Output Shape:	None, 75, 300
Dropout	Output Shape:	None, 75, 100
Bidirectional (LSTM)	Output Shape:	None, 75, 300
Dropout	Output Shape:	None, 75, 100
Timedistributed (dense)	Output Shape:	None, 75, 300

Training our own Cyber Entity Extractor

Training Data

Feature Extraction

Architecture

Assessments

Assessment

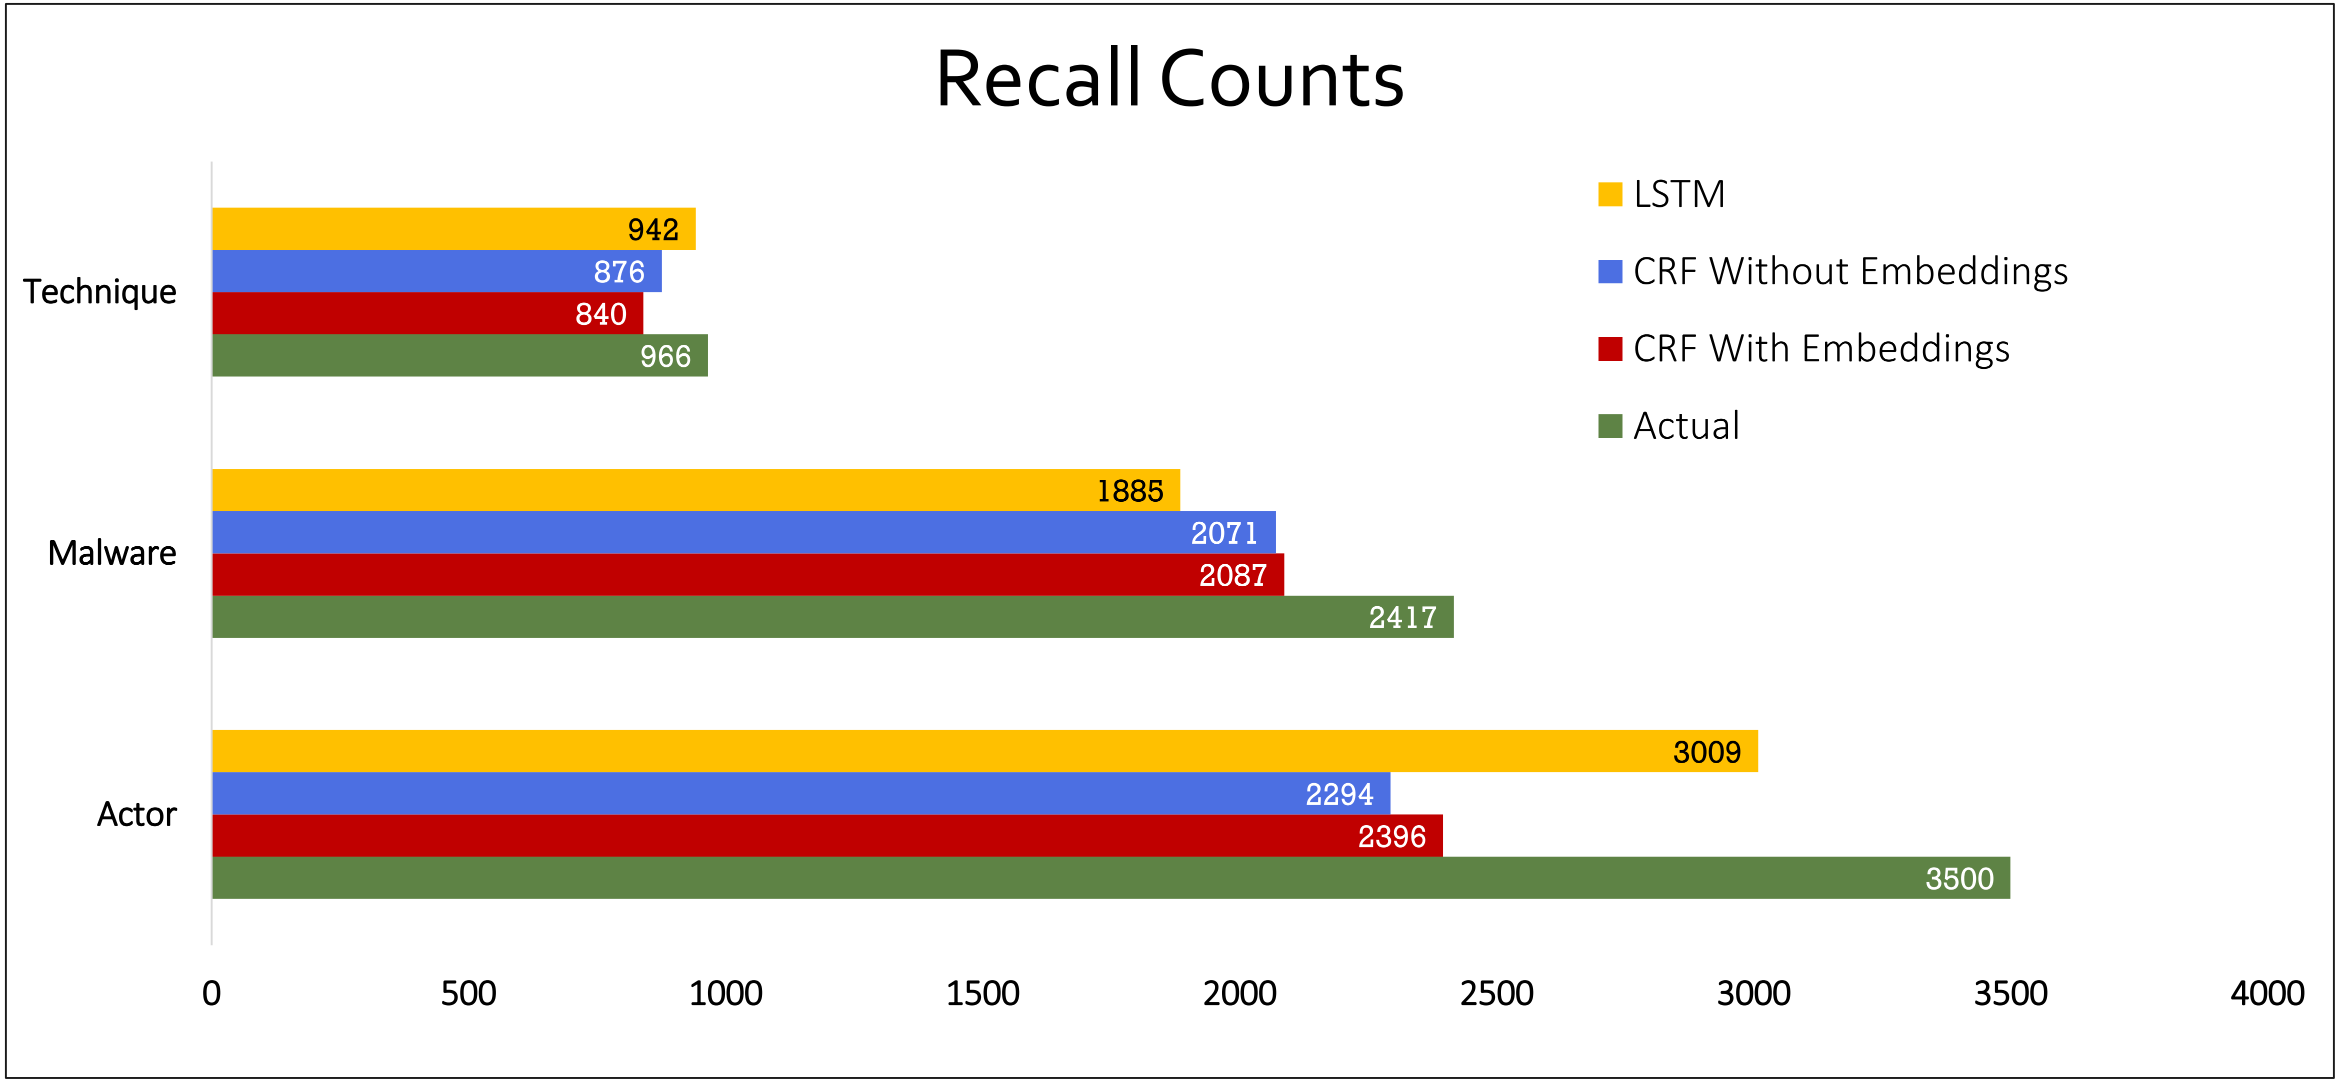
Recall Counts
LSTM
CRF Without Embeddings
CRF With Embeddings
Actual
Technique
942
876
840
966
Malware
1885
2071
2087
2417
Actor
3009
2294
2396
3500
0
500
1000
1500
2000
2500
3000
3500
4000

Assessment

Precision = TP/(FP+TP)
Recall = TP/(TP+FN)

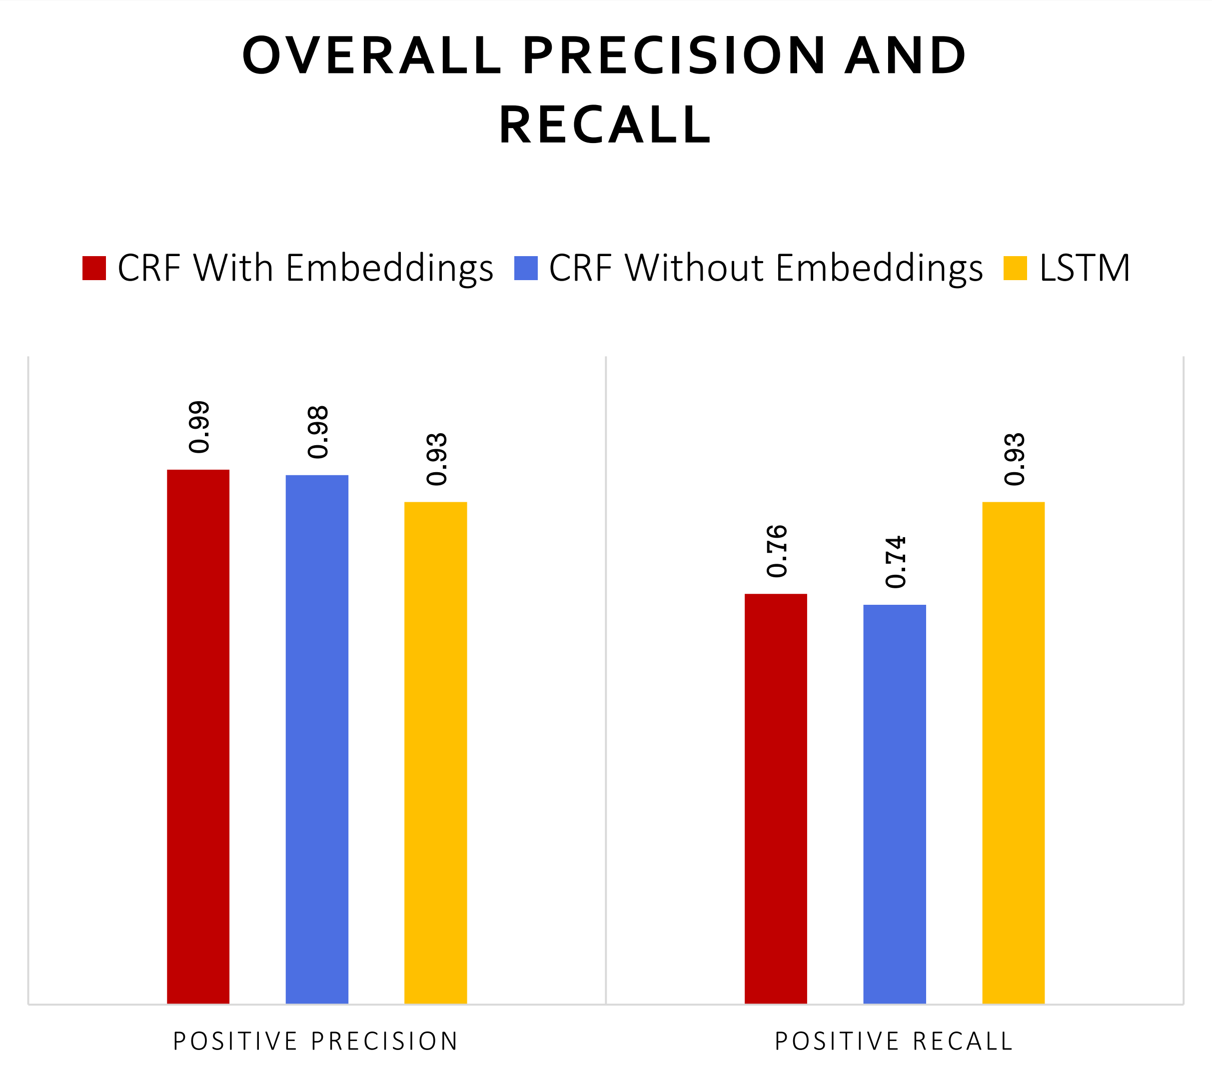

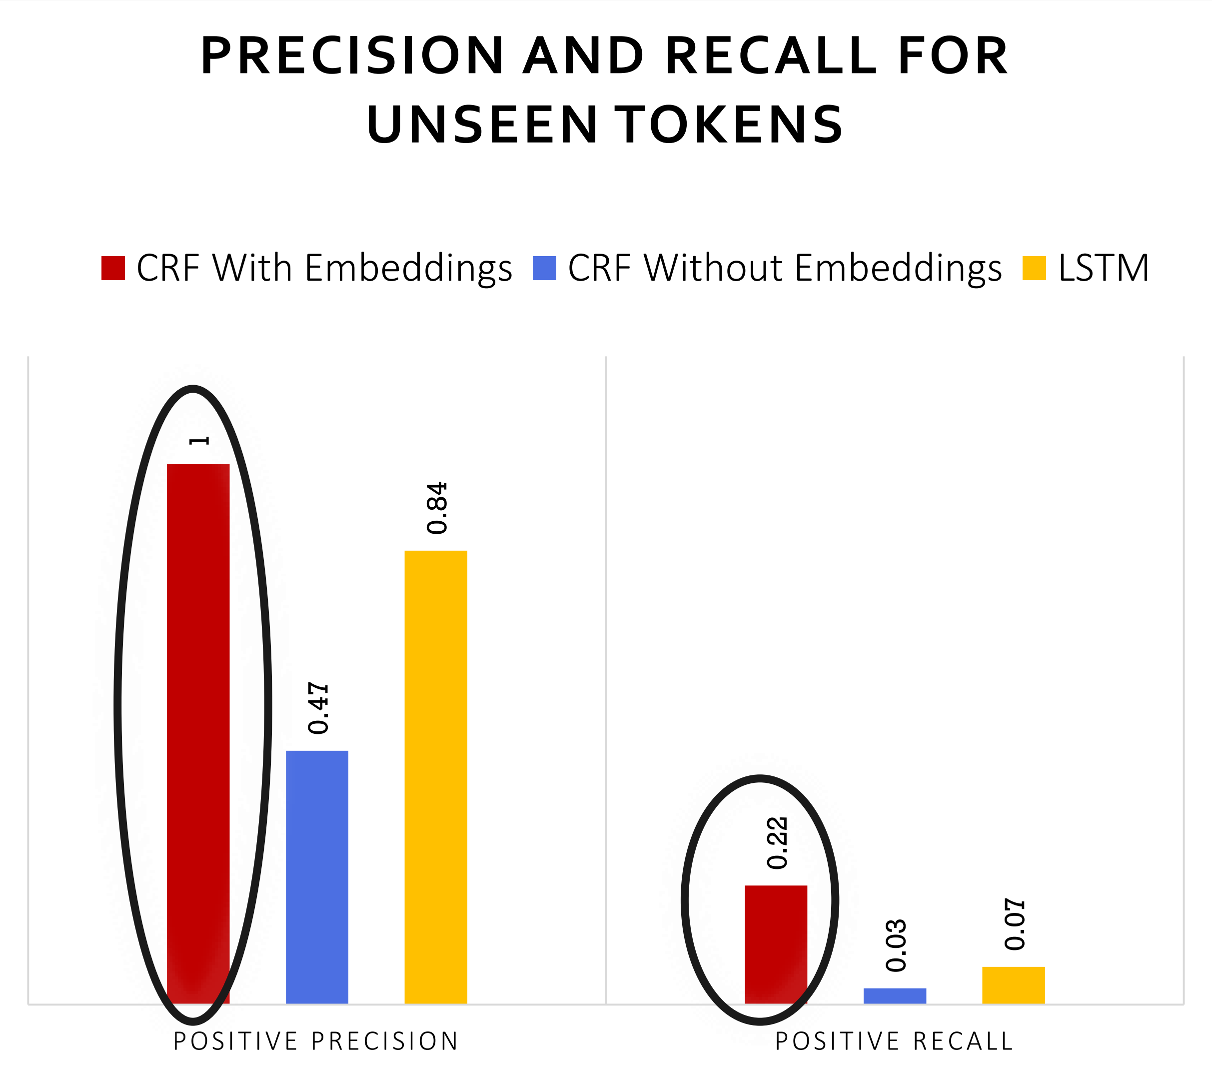

Demo

Custom Entity Extractio ×

localhost:8080/

Custom Entity Extraction for Threat Intelligence

A demonstration of using machine learning to extract malware classification entities from security publications.

Model: CRF With Embeddings

Enter malware article text.

0/2000

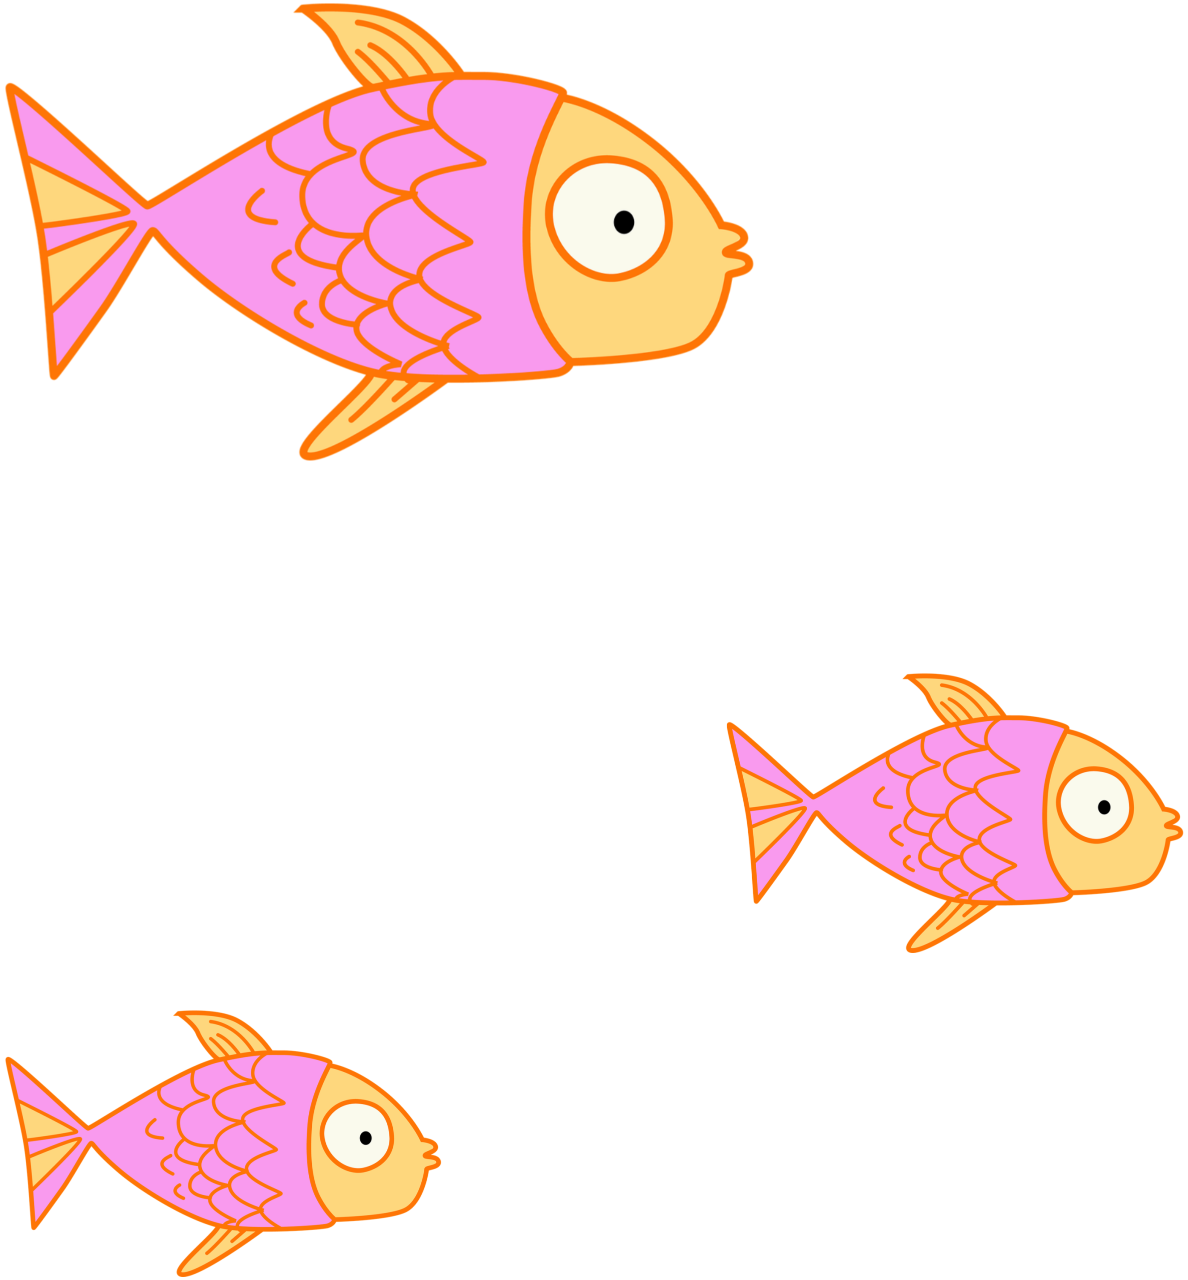

Next Steps…

- Attention networks
- Data Augmentation
- Sophisticated Relationship Extraction
- Temporal relationships

Driving Impact

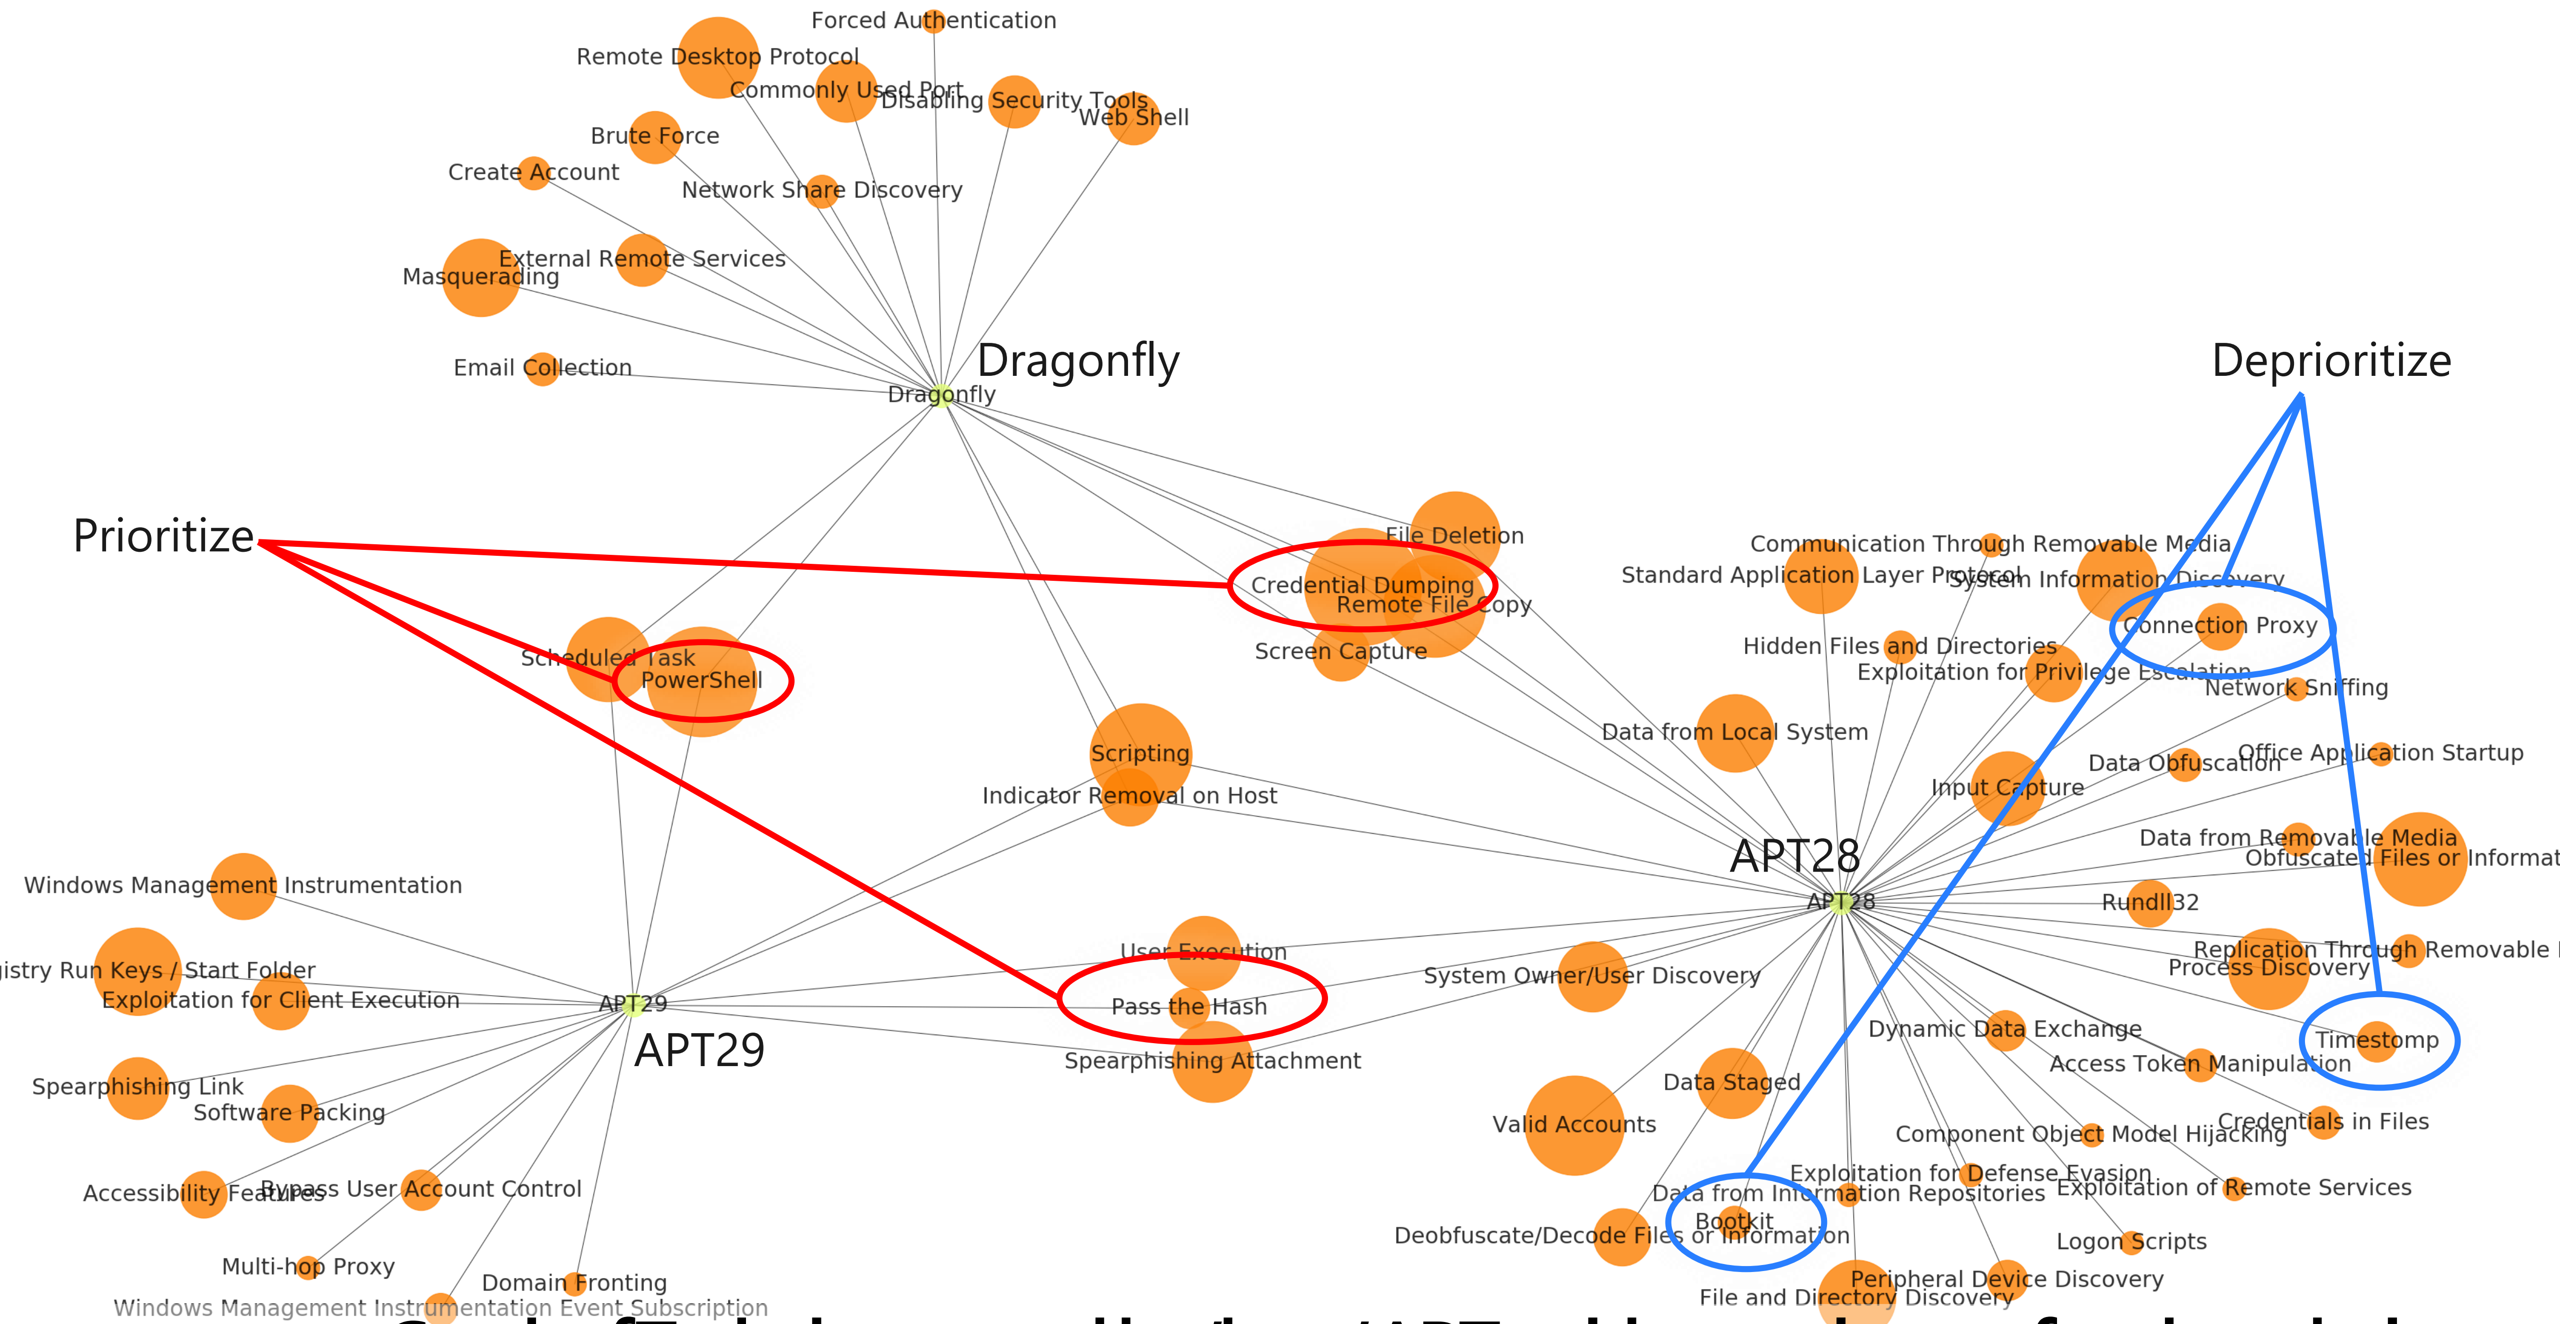

Graph of Techniques used by /bear/ APTs with prevalence of each technique

https://mitre-attack.github.io/caret/#/

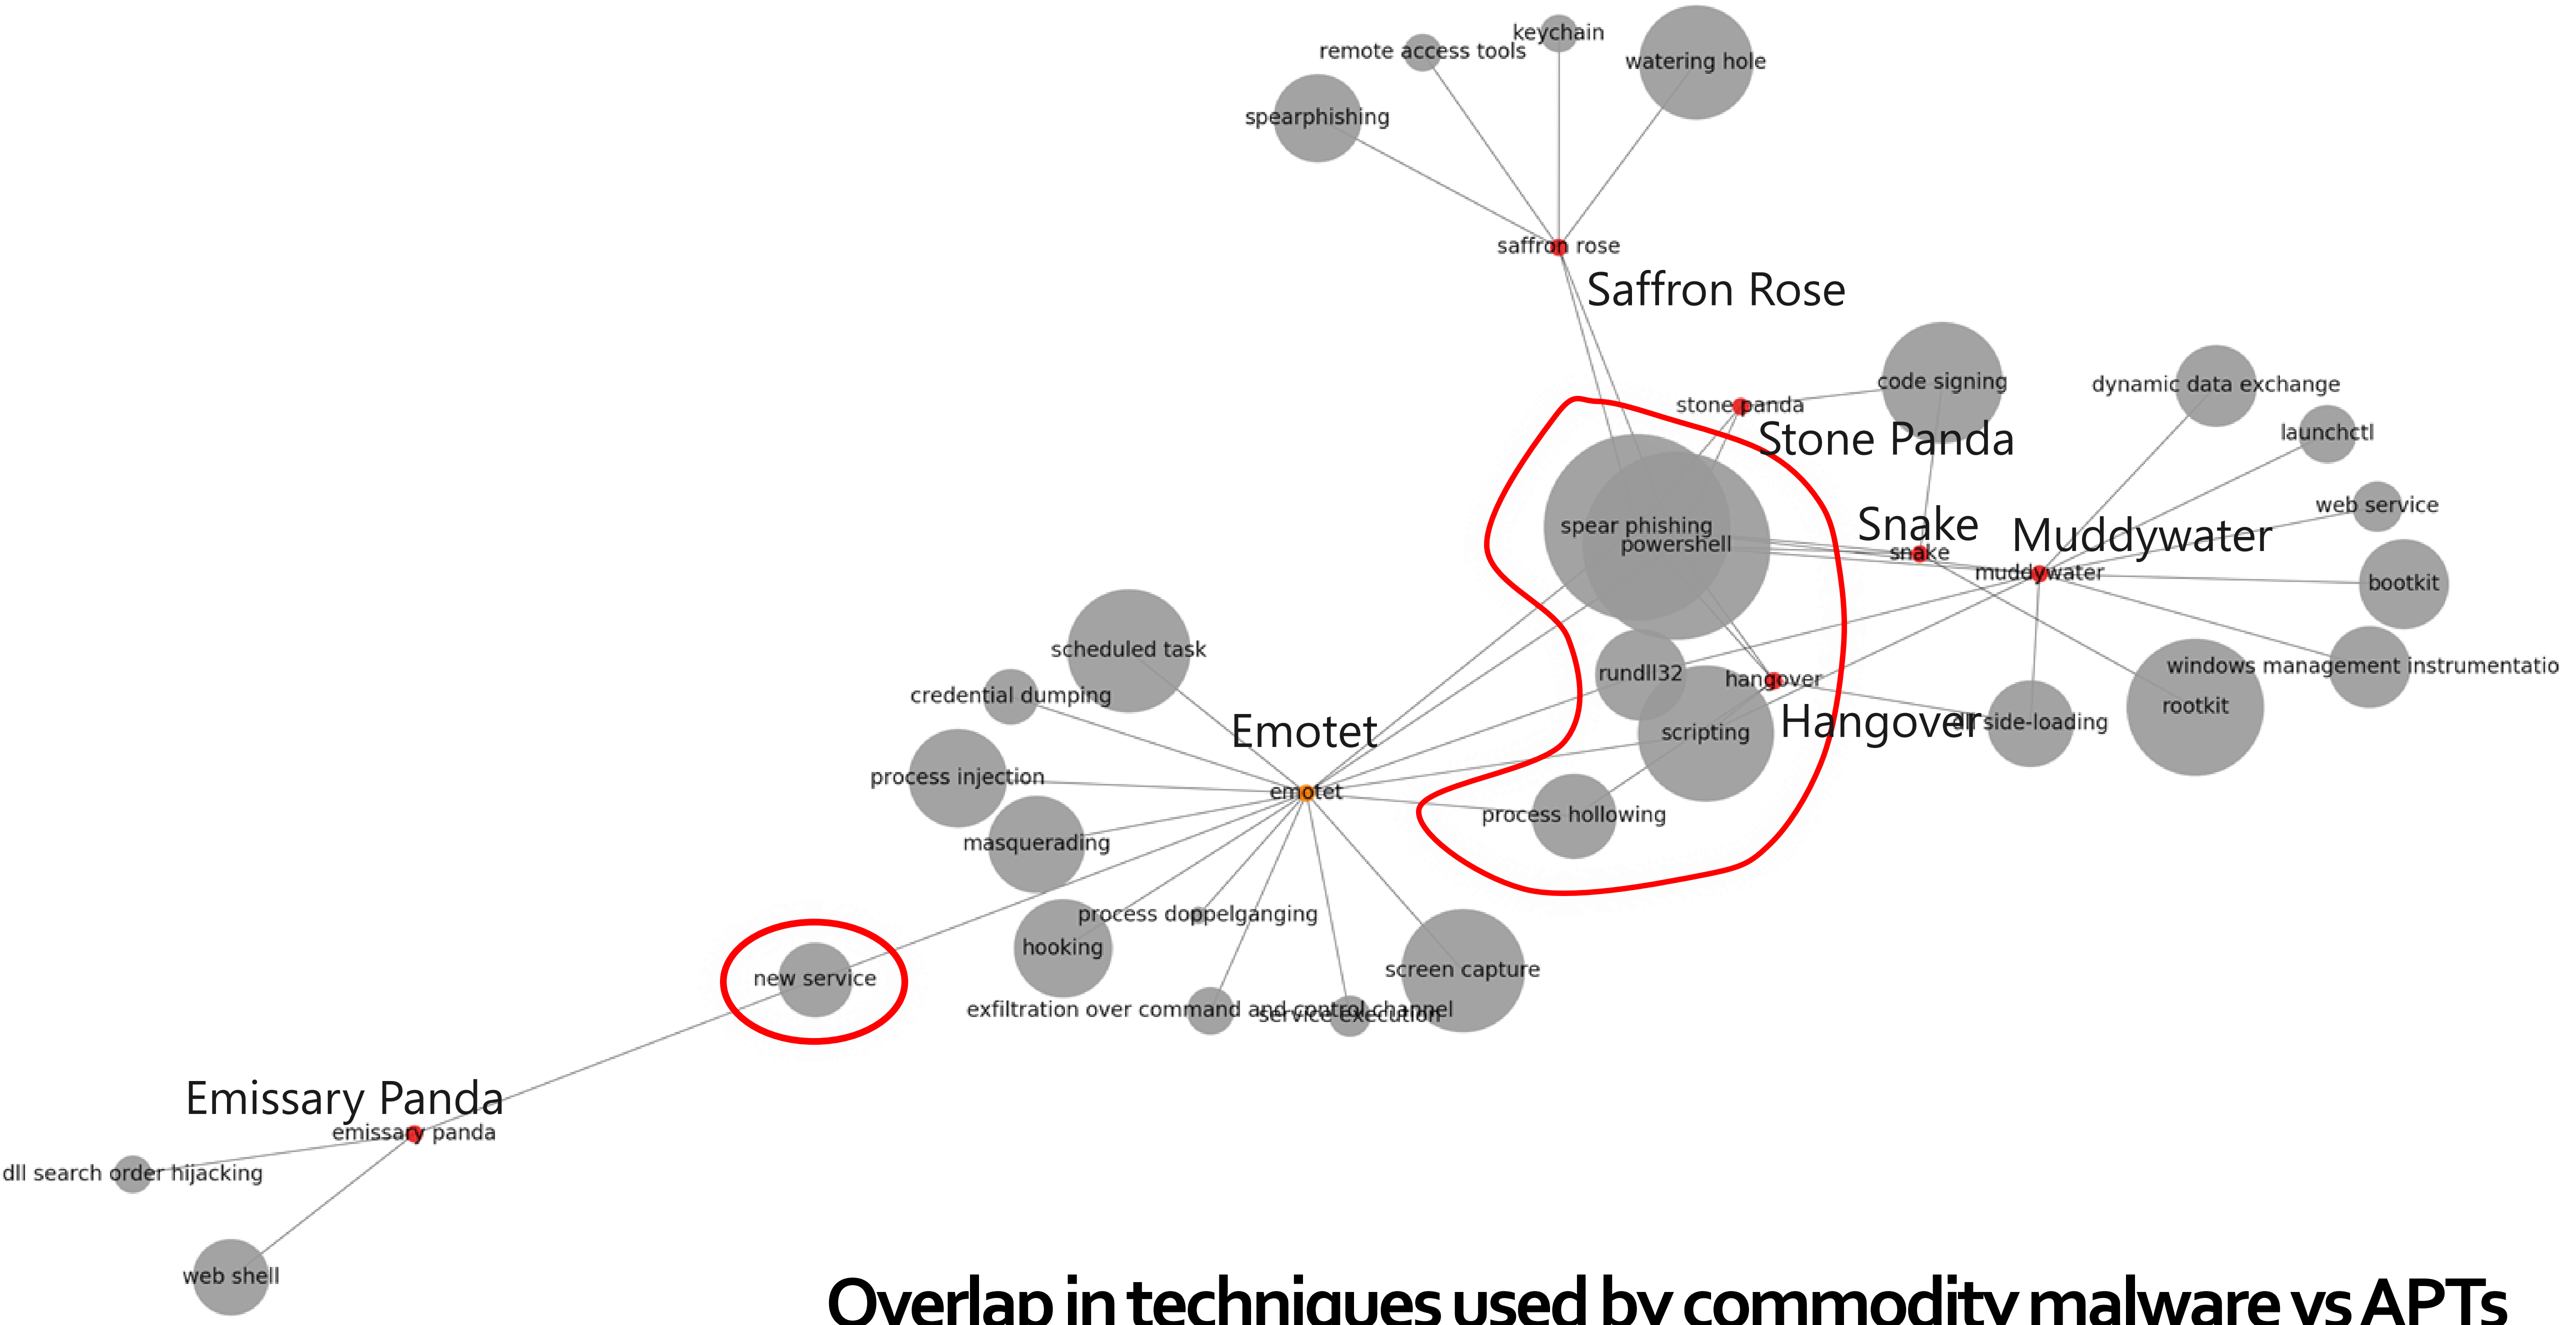

Overlap in techniques used by commodity malware vs APTs

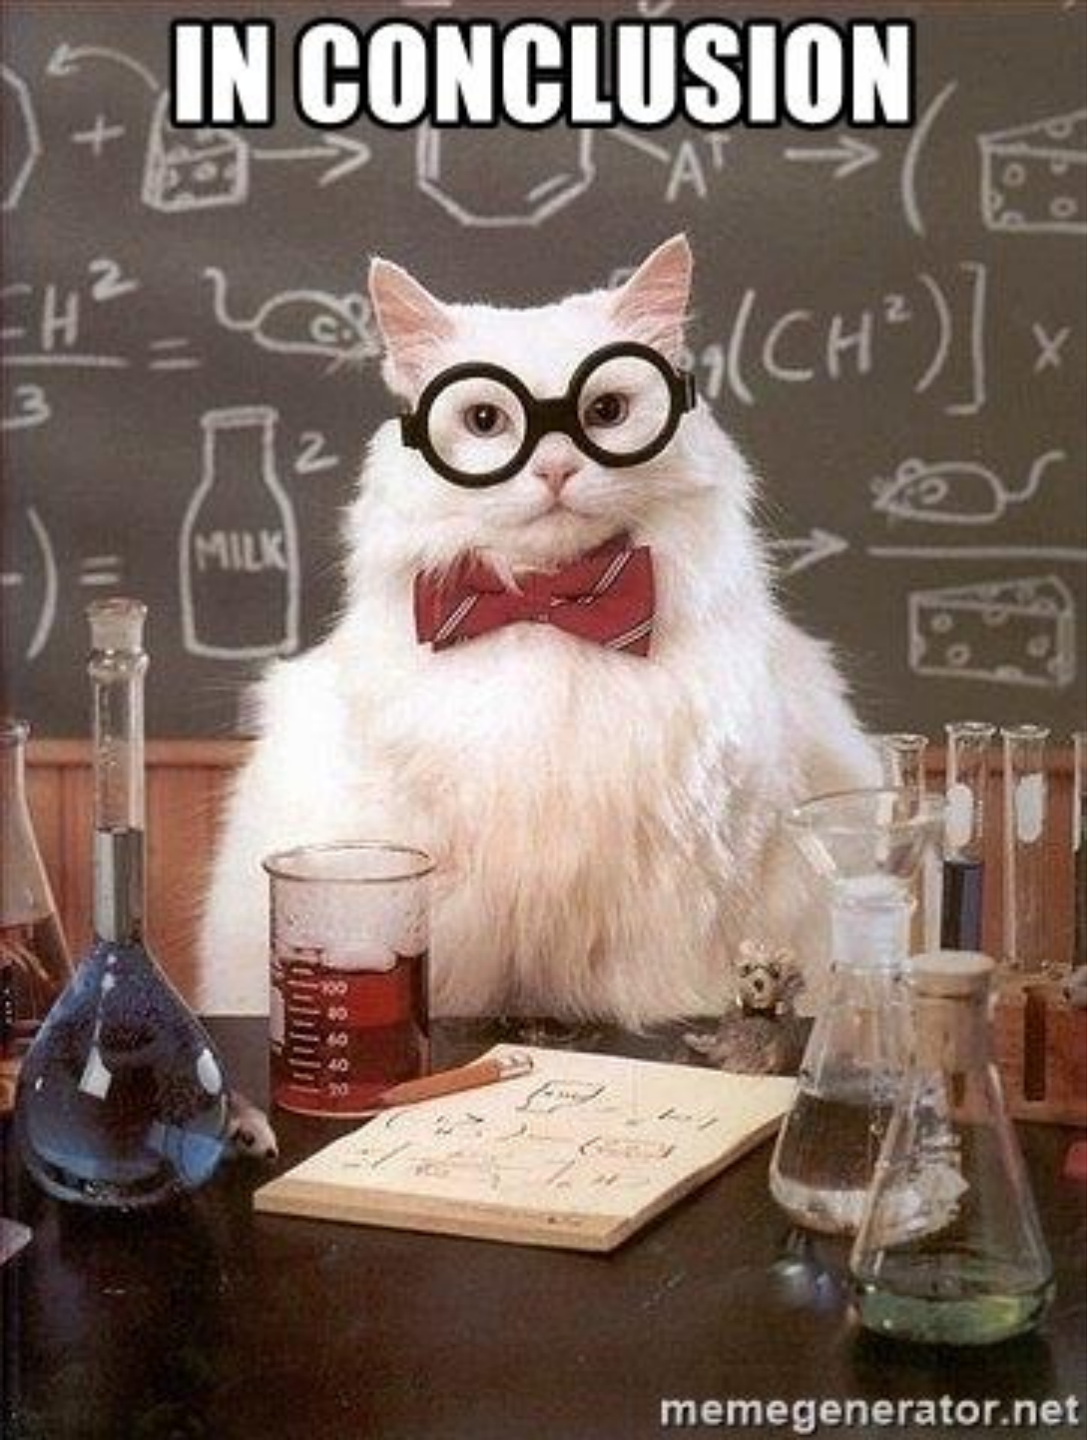

- Move beyond IOC feeds
- Rich unstructured data can be extracted with Machine Learning
    - Graphs
    - Timelines
- We can use this to make better decisions to improve security of our orgs

Acks/Q&A/Thanks

- Contributors:
  - Arun Gururajan, Daewoo Chong and Jugal Parikh for Data Science Expertise
  - Peter Cap and Jessica Payne for Threat Intelligence Expertise
  - Chris Ackerman for the demo website
  - Karen Lavi for encouragement, better presentation

bhavna.soman@microsoft.com
@bsoman3