




SKYLIGHT
ZERO COMPROMISE



Attacking Machine Learning The Cylance Case Study

BSides Sydney 2019

About Us

Not even AI power users...



Adi Ashkenazy
CEO, Skylight Cyber

Wears T-Shirts in Corporate Headshots

[yes, I am wearing it now too!]

Heavy on the offensive cyber side (Government)

Red-team automation



Shahar Zini
CTO, Skylight Cyber

Category 5 stage fright

Sydney based consultancy

Help companies navigate cyber security



In a nutshell

Why is this important?

What are we looking to achieve?

AI in Cyber for people who understand quickly

How we approached the problem and reversing the product

Results!

Publication and Feedback

Questions

Challenge



Silver Bullet
Hunting

Assess
Technology

“The product is as close as you can get to a silver bullet in our space. Greater than 99% efficacy and protection against nearly every zero-day malware

WHY CYLANCE?

AI Centric, can buy it off the shelf, consistently ranks high

Their marketing didn't help!



Zero-Day Attacks

AI model prevents zero-day payloads from executing

Forrester Report: Cylance Provides 251% ROI

Classification with AI/ML - The Basics

“In machine learning and statistics, classification is the problem of identifying to which of a set of categories (sub-populations) a new observation belongs, on the basis of a training set of data containing observations (or instances) whose category membership is known”

Wikipedia



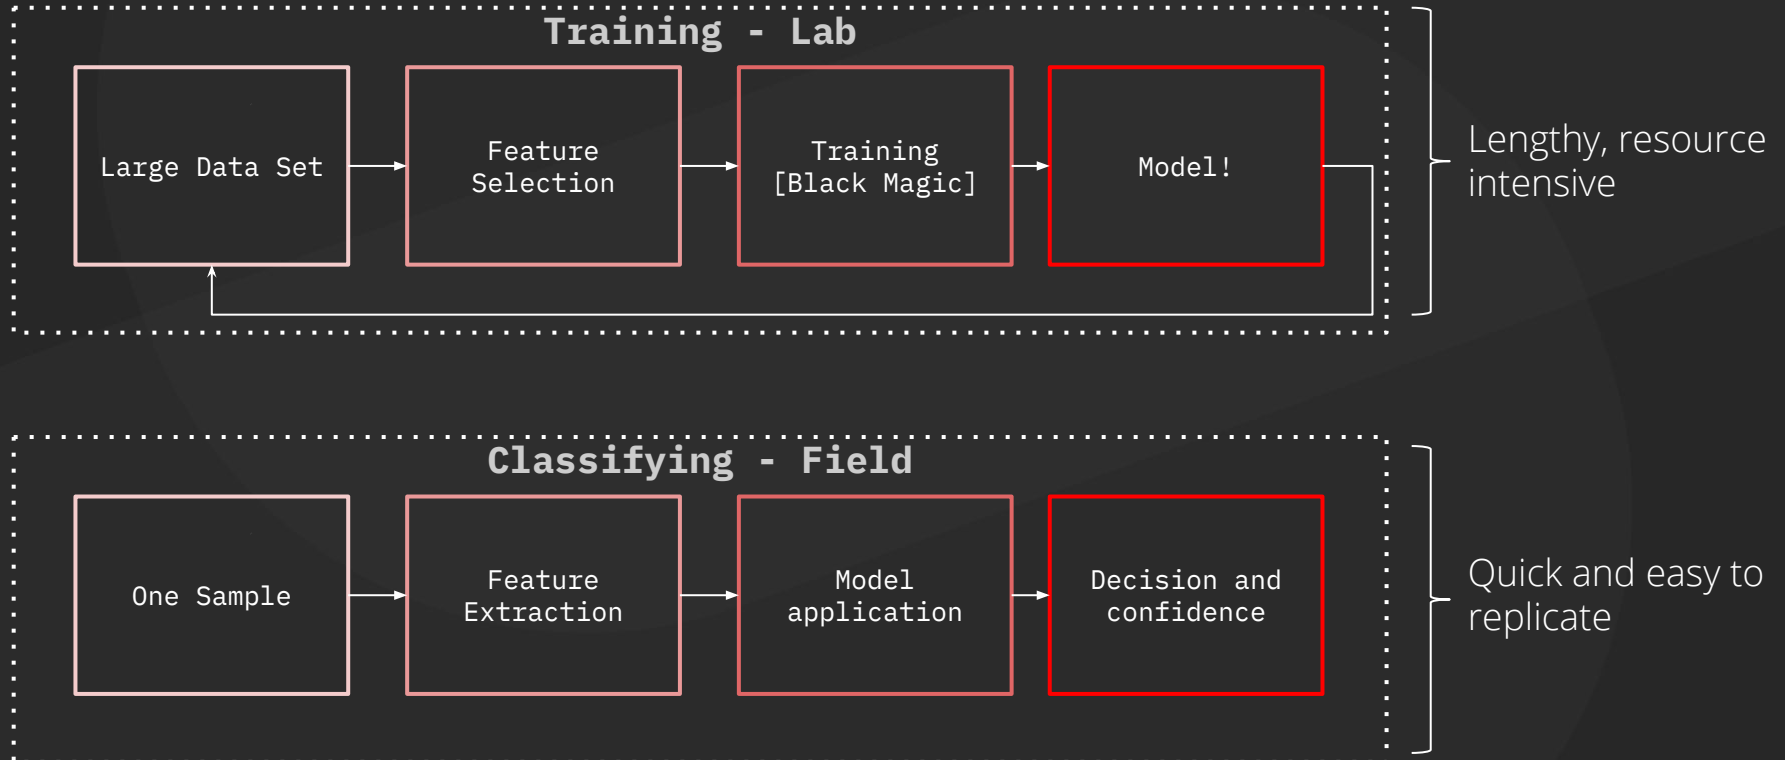
What is THAT?

Cat

Dog



Classification with AI/ML




An offensive mindset

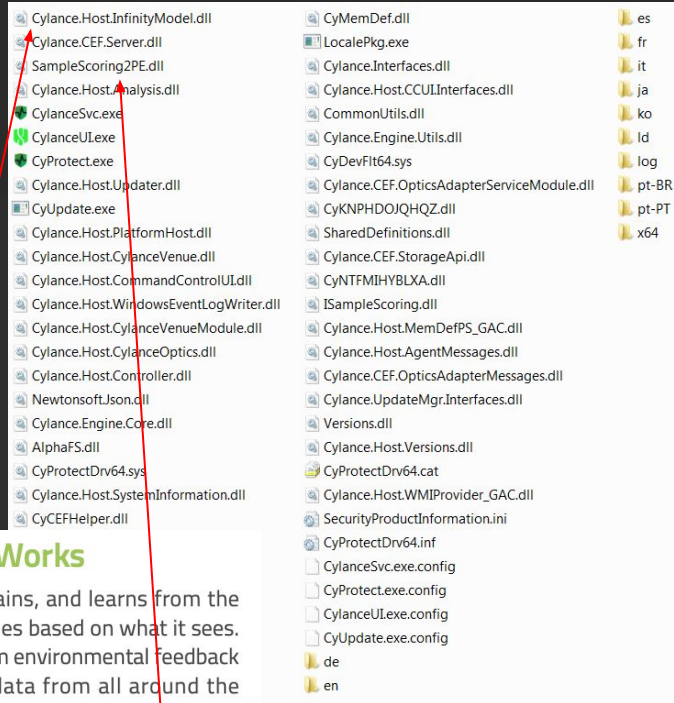
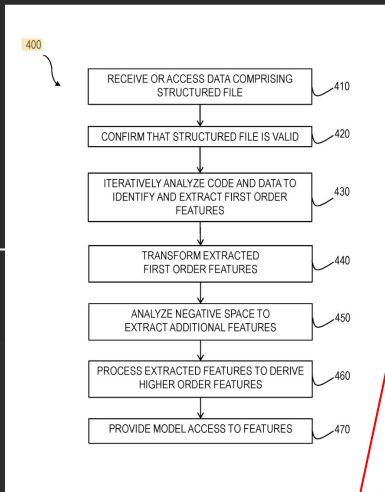
- Classification is innately naive
- A model is only as good as its data
- How would we fool the bird vs. human classifier?



The OpSec Paradox



OSINT



Marketing & Legal

Patents

White Papers & Booklets

Conference Talks

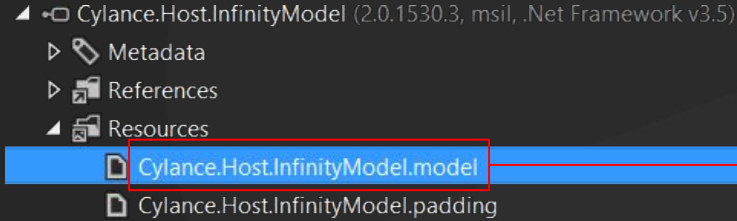


How CylanceINFINITY Works


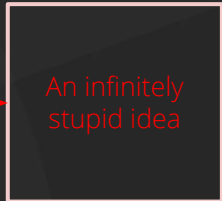

CylanceINFINITY collects data, trains, and learns from the data, then calculates likely outcomes based on what it sees. It's constantly getting smarter from environmental feedback and a constant stream of new data from all around the world.

To achieve its magic, CylanceINFINITY first collects vast amounts of data from every conceivable source. Next, it extracts features that we have defined to be uniquely atomic characteristics of the file depending on its type (.exe, .dll, .com, .pdf, .java, .doc, .xls, .ppt, etc.). Then, it constantly adjusts to the real-time threatscape and trains the machine learning system to make better decisions. Finally, for each query to CylanceINFINITY, we classify the data as good or bad.

Extracting the Model



```
17
18 namespace Cylance.Engine.Core.Ensemble
19 {
20     public class EnsembleReader : ILogAccess, IEnsembleHeader, IDisposable
21     {
22         protected const int RandomHeaderSize = 3072;
23         protected const string KeyAndIv = "I am decrypting Cylance's intellectual property";
24         protected Stream _stream;
25         protected byte[] _activeKey;
26         protected byte[] _activeIV;
27         protected bool _loadSectionData;
```



Our own classifier

Engineering Masterpiece!

Let's build our own classifier so we can dynamically debug and follow the code


```
static void Main(string[] args)
{
    SampleScoreFactory2PE factory = new SampleScoreFactory2PE();
    SampleScoring2PE scorer = factory.Create("test_model.bin") as SampleScoring2PE;
    Stream test_file = File.Open("mimikatz_with_slight_modification.exe", FileMode.Open);
    Dictionary<string, object> extraData;
    double score = scorer.ComputeScore(test_file, out extraData);
}
```

Watch 1

| Name | Value |
|-------|----------------------|
| score | -0.85276468809127071 |

LocalAnalyzeItem, C:\Users\Administrator\Desktop\mimikatz_with_slight_modification.exe score -852 detector execution_control

Anti-Tampering & Obfuscation



```
[assembly: CompilationRelaxations(8)]  
[assembly: RuntimeCompatibility(WrapNonExceptionThrows = true)]  
[assembly: Debuggable(/*Could not decode attribute arguments.*/) ]  
[assembly: AssemblyTitle("SampleScoring2 (PE)")]  
[assembly: AssemblyConfiguration("")]  
[assembly: AssemblyCompany("Cylance, Inc.")]  
[assembly: AssemblyProduct("SampleScoring")]  
[assembly: AssemblyCopyright("Copyright © Cylance, Inc.")]  
[assembly: AssemblyTrademark("")]  
[assembly: ComVisible(false)]  
[assembly: Guid("EAA9899F-E9C9-439F-83FB-3045CB68E5C")]  
[assembly: AssemblyFileVersion("4.3.2.3244")]  
[assembly: AssemblyInformationalVersion("1.0.0")]  
[assembly: PoweredBy("Powered by SmartAssembly 6.9.0.114")]  
[assembly: SuppressIldasm]  
[assembly: AssemblyVersion("4.3.2.0")]  
[module: UnverifiableCode]
```

USELESS


| | | | | | |
|---------|-------------|-------------|-------------|-------------|-------------------|
| DA9F0h: | 20 41 73 73 | 65 6D 62 6C | 79 52 65 66 | 73 00 00 00 | AssemblyRefs... |
| DA00h: | 3C 4D 6F 64 | 75 6C 65 3E | 00 00 00 00 | 53 79 73 74 | <Module>...Syst |
| DA10h: | 65 6D 2E 52 | 75 6E 74 69 | 6D 65 2E 43 | 6F 6D 70 69 | em.Runtime_Compil |
| DA20h: | 6C 65 72 53 | 65 72 76 69 | 63 65 73 2E | 53 75 70 70 | lerServices_Supp |
| DA30h: | 72 65 73 73 | 49 6C 64 61 | 73 6D 41 74 | 74 72 69 62 | ressIldasmAttrib |
| DA40h: | 75 74 65 00 | 6D 00 73 00 | 63 00 6F 00 | 72 00 6C 00 | lter.m.s.c.o.r.l. |
| DA50h: | 69 00 62 00 | 00 00 00 00 | 25 73 20 25 | 73 20 25 73 | i.b...%s %s %s |
| DA60h: | 00 00 00 00 | 2F 2F 20 43 | 6C 61 73 73 | 65 73 20 64 | ...// Classes d |

```
122 if ((long) this.int_10 != (long) this.uint_1)  
123 {  
124     this.uint_1 = (uint) this.int_10;  
125     this.byte_3 = new byte[(int) this.uint_1];  
126 }  
127 this.byte_3[(int) this.uint_1 - 1] = (byte) 0;  
128 for (this.int_8 = 5; this.int_8 > 0; this.int_8 = this.int_8 - 1)  
129 {  
130     if (this.uint_2 >= 10)  
131     {  
132         int num8 = (int) this.uint_13 << 8;  
133         byte[] byte0_2 = this.byte_0;  
134         int int15_2 = this.int_15;  
135         this.int_15 = int15_2 + 1;  
136         int index2 = int15_2;  
137         int num9 = (int) byte0_2[index2];  
138         this.uint_13 = (uint) (num8 | num9);  
139         this.uint_2 = this.uint_2 - 10;  
140     }  
141     else  
142     {  
143         this.int_0 = ...  
144         return this.memoryStream_0.ToArray();  
145     }  
146 }  
147 uint uint14_2 = this.uint_14;  
148 this.uint_14 = uint14_2 + 16;  
149 if (uint14_2 > 335544320)  
150     goto label_138;  
151     else  
152         goto case 2;  
153 }  
154     else  
155         goto label_137;  
156 }  
157 }  
158     goto label_136;  
159 }  
160 case 2:  
161     this.int_7 = (int) this.uint_4 & (int) this.uint_9;  
162     this.int_5 = (this.int_12 << 4) + this.int_7;  
163     this.int_1 = 6;  
164     uint uint14_3 = this.uint_14;  
165     this.uint_14 = uint14_3 + 10;  
166     if (uint14_3 < 335544320)  
167         goto case 4;
```


WARNING

Parsing the Properties

```
this.Observations.Add("AWImportCount", (object) (class63.class56_0.int_0 + class63.class56_0.int_1));  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("TransactedImportCount", (object) class63.class56_0.int_2);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("FileImportCount", (object) class63.class56_0.int_3[0]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("RegistryImportCount", (object) class63.class56_0.int_3[1]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("DebuggeeImportCount", (object) class63.class56_0.int_3[2]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("DebuggerImportCount", (object) class63.class56_0.int_3[3]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("DestructiveImportCount", (object) class63.class56_0.int_3[4]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("DirectExecutableImportCount", (object) class63.class56_0.int_3[5]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("EscapedImportCount", (object) class63.class56_0.int_3[6]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("ExtensionBaseImportCount", (object) class63.class56_0.int_3[7]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("FilePersistenceImportCount", (object) class63.class56_0.int_3[8]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("InjectionImportCount", (object) class63.class56_0.int_3[9]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("InternetImportCount", (object) class63.class56_0.int_3[10]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("IPCImportCount", (object) class63.class56_0.int_3[11]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("ProcessEnumerationImportCount", (object) class63.class56_0.int_3[12]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("RemoteAudioVideoImportCount", (object) class63.class56_0.int_3[13]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("RemoteInputImportCount", (object) class63.class56_0.int_3[14]);  
// ISSUE: reference to a compiler-generated field  
this.Observations.Add("UserEnumerationImportCount", (object) class63.class56_0.int_3[15]);
```



Parser



| Property | Value |
|----------------|-------------|
| Linker version | 5.1 |
| Num sections | 5 |
| Section casing | Uppercase |
| Entropy | 0.2315 |
| Timestamp | 13102382120 |


⋮

| | |
|------------------|-------|
| Max section size | 827Kb |
| CLR version | 4.0 |
| #UI imports | 98 |
| #Process imports | 14 |
| #imports | 412 |

Building the Feature Vector



Property->Feature
Converter

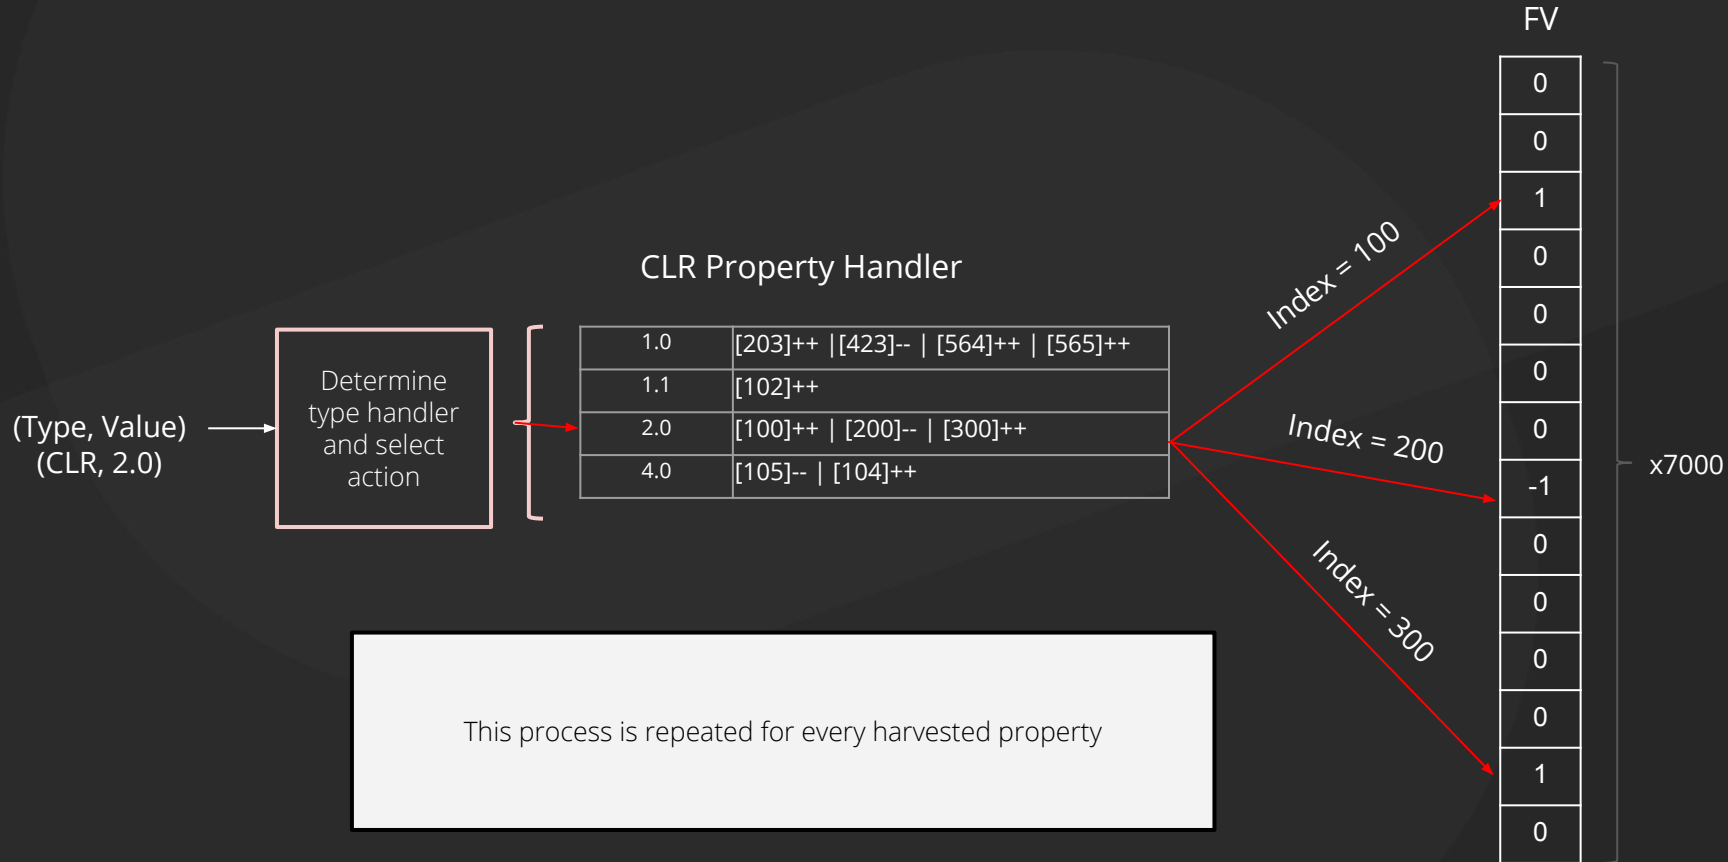


| | | | | | | | | | | | | | | | | | | | |
|------------------|---|----------------|------------------------|----------------|------------------------|----------------|------------------------|----------|--|----------|--|----------|--|----------|--|----------|--|----------|--|
| Property Type #1 | <table border="1"><tr><td>Value #1</td><td></td></tr><tr><td>Value #2</td><td></td></tr><tr><td>Value #3</td><td></td></tr><tr><td>Value #4</td><td></td></tr><tr><td>Value #5</td><td></td></tr><tr><td>Value #6</td><td></td></tr><tr><td>Value #7</td><td></td></tr><tr><td>Value #8</td><td></td></tr><tr><td>Value #9</td><td></td></tr></table> | Value #1 | | Value #2 | | Value #3 | | Value #4 | | Value #5 | | Value #6 | | Value #7 | | Value #8 | | Value #9 | |
| Value #1 | | | | | | | | | | | | | | | | | | | |
| Value #2 | | | | | | | | | | | | | | | | | | | |
| Value #3 | | | | | | | | | | | | | | | | | | | |
| Value #4 | | | | | | | | | | | | | | | | | | | |
| Value #5 | | | | | | | | | | | | | | | | | | | |
| Value #6 | | | | | | | | | | | | | | | | | | | |
| Value #7 | | | | | | | | | | | | | | | | | | | |
| Value #8 | | | | | | | | | | | | | | | | | | | |
| Value #9 | | | | | | | | | | | | | | | | | | | |
| Property Type #2 | <table border="1"><tr><td>Value Range #1</td><td>Sequence of actions #1</td></tr><tr><td>Value Range #2</td><td>Sequence of actions #2</td></tr><tr><td>Value Range #3</td><td>Sequence of actions #3</td></tr></table> | Value Range #1 | Sequence of actions #1 | Value Range #2 | Sequence of actions #2 | Value Range #3 | Sequence of actions #3 | | | | | | | | | | | | |
| Value Range #1 | Sequence of actions #1 | | | | | | | | | | | | | | | | | | |
| Value Range #2 | Sequence of actions #2 | | | | | | | | | | | | | | | | | | |
| Value Range #3 | Sequence of actions #3 | | | | | | | | | | | | | | | | | | |
| | ⋮ | | | | | | | | | | | | | | | | | | |
| Property Type #N | <table border="1"><tr><td>Value #1</td><td></td></tr><tr><td>Value #2</td><td></td></tr><tr><td>Value #3</td><td></td></tr><tr><td>Value #4</td><td></td></tr><tr><td>Value #5</td><td></td></tr></table> | Value #1 | | Value #2 | | Value #3 | | Value #4 | | Value #5 | | | | | | | | | |
| Value #1 | | | | | | | | | | | | | | | | | | | |
| Value #2 | | | | | | | | | | | | | | | | | | | |
| Value #3 | | | | | | | | | | | | | | | | | | | |
| Value #4 | | | | | | | | | | | | | | | | | | | |
| Value #5 | | | | | | | | | | | | | | | | | | | |


Property Type
Handler



Building the Feature Vector



Linear Algebra, How I Missed You

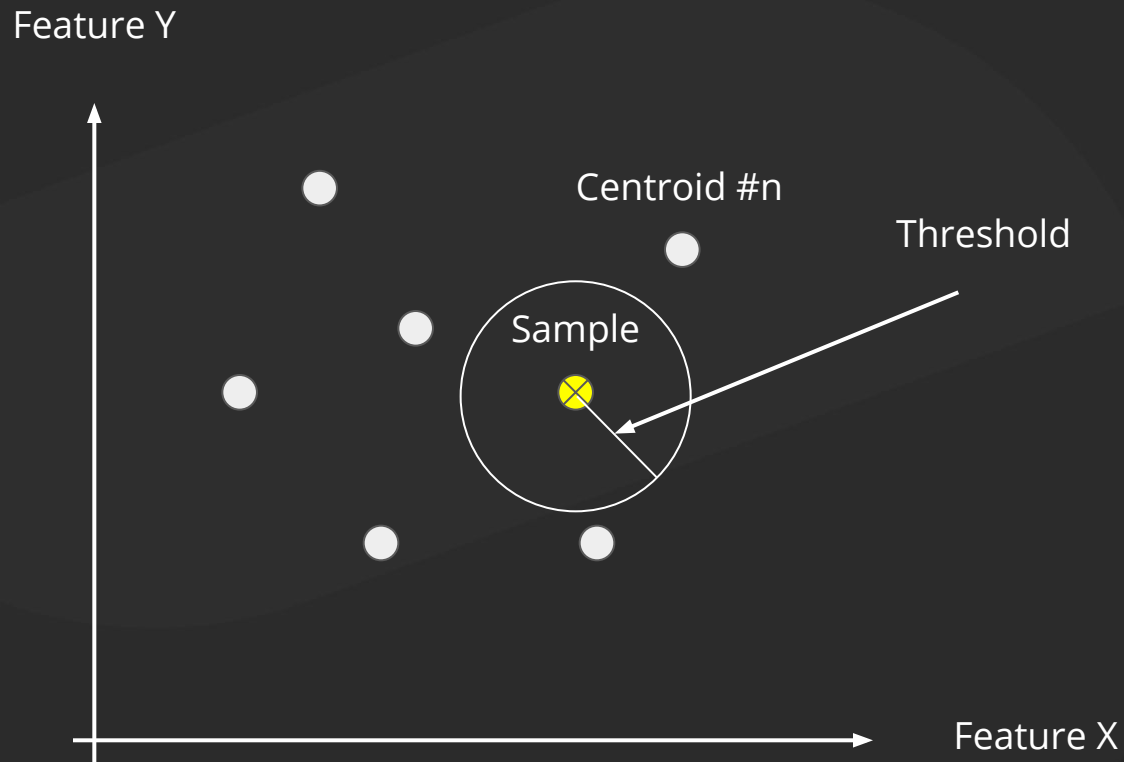



Populated by
property
processing

The Model Core

[-1, 1]
Finally!

White/Black-Listing



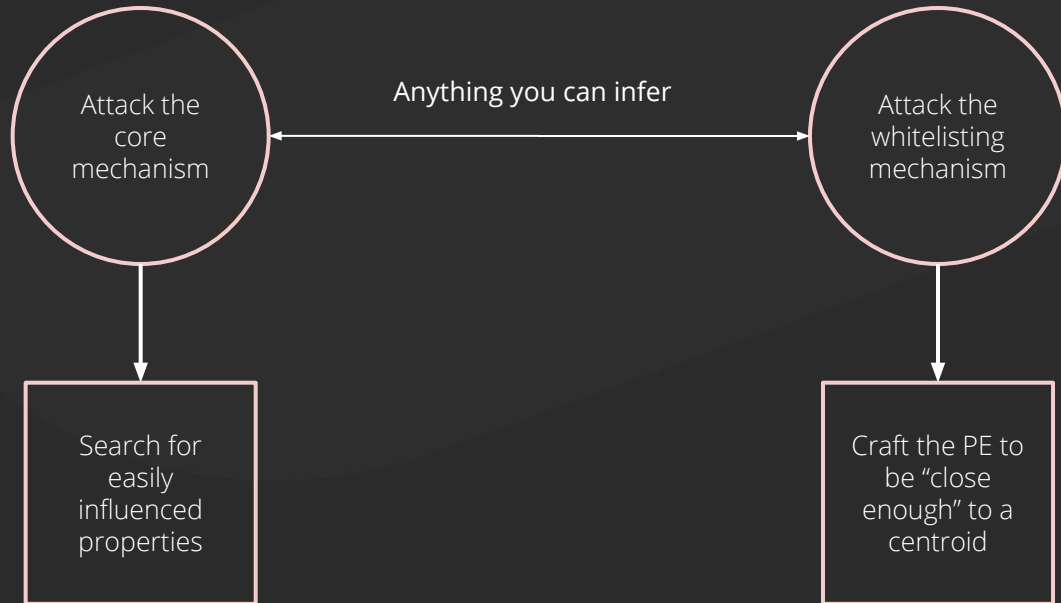
Rocket What?

Watch 1

| Name | Value |
|---|--|
| [[white, System.Collections.Generic.Dictionary`2[System.String,Cylance.Engine.Core.Centroids.Centroid]] | |
| "white" | |
| Count = 0x00000010 | |
| | [[fullspace_dalek_white_trusted_kbhomes_preso_centroid_1-C1CEA4C9136C7909AA1655C4566B21F948514666D90289334A550DDC5DEE56B6, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_torqtek_packer_pcgaurd_centroid_2-6A6B21AAF6B8714406E3E6FA238F51EAACC2C8F8AC6FEA3AAAA90F0352626F28, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_ztp_centroid_3-BDE4ED001B100FB08FC8C32CFAE4BC181EC4AFB3E7DD464E92ADDEAAF5DA2B9F, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_cylance_inspect_centroid_0-01CF56B85EBF820765D38EFD9C9C6677ECF4C08F72E55AEE9B5899662A79F078B, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_torqtek_packer_pcgaurd_centroid_3-5EBF6E2EC35DA09AA1D9A255AE790660B6050170C68F293ED8019D3661CA8138, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_usmt_centroid_0-C16D6ED90F3CC7F4D1014CA1CFE4FF202EEFF739ECE647CE79B17F7237AB9CF6, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_ztp_centroid_2-F31D416F55A43ED4B273CD04B0279148E5425C0F88C29BA1AD6745895DEE6E3A, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_torqtek_packer_pcgaurd_centroid_4-3EAC4DF016685F6A8D3DE0D09A47DCBC144166505E456E7E91EEC8964158856D, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_ztp_centroid_1-69270CBF3FFCB5FBA3803C84835971C105020BA83B9848A32EF64EB61EF734BF, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_rockleaguel_centroid_0-E3D3D8B2893F36C10A71B203A064FDA3E8400A6D84095C4024A1241234C8E01B, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_ztp_centroid_4-94A72D94AE2AE5AA9BF1830C725E7977AF43F06C7AD06A7F2F60F128FF45887, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_torqtek_packer_pcgaurd_centroid_0-D7EFBEE08E2E6B715F1D25FFB4BC55E3D40A89B5ECD97676273E77C84DA2BC84, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_torqtek_packer_pcgaurd_centroid_1-33B621DF787852D99C34F3B9ACC7CA10E01AD8E55B740CF1AE7805223A3CCB26, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_kbhomes_preso_centroid_0-69C6C38C3A1BC4DABC0E2344F377EA3D3B4E10BC36CF4366AEC88BDAA3FD40898, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_ztp_centroid_5-7FC4B110F3ED9B7945DC3EC309265115AF41BCCE2FB341D1F598E64DE35D4CE9, Cylance.Engine.Core.Centroids.Centroid]] |
| | [[fullspace_dalek_white_trusted_ztp_centroid_0-9CE022C5886014AA4C5B13CD1F6F241D561A7889881B830B1ABB356F692D7B11, Cylance.Engine.Core.Centroids.Centroid]] |

Hmmm... This could be interesting,
hold that thought

Let's Pause and Hypothesise



Strings Galore

Location of the String Type handler

WOAH, that's a large handler!

```
for (int index = 0; index < this.imagePEFile_0.Strings.Length; ++index)
{
    if (!this.method_26(this.imagePEFile_0.Strings[index].S, 95088, 854069, 0))
        this.method_14(this.list_0[0], 15166118410741992125UL, 2847678, 0);
}
```

Process property function

Strings Galore, Contd.

| | |
|--------------|------------------------------------|
| HASH(Str #1) | [203]++ [423]-- [564]++ [565]++ |
| HASH(Str #2) | [1020]++ |
| HASH(Str #3) | [866]++ [533]-- |
| HASH(Str #4) | [53]++ [4]-- [2464]++ [5432]++ |
| HASH(Str #5) | [4500]++ [3223]-- |
| HASH(Str #6) | [10]++ [400]-- [3444]++ |
| HASH(Str #7) | [453]++ |

⋮

| | |
|-------------------|-----------------------------------|
| HASH(Str #854063) | [23]++ [25]-- [55]++ |
| HASH(Str #854064) | [6088]++ [48]-- [4332]++ [2]++ |
| HASH(Str #854065) | [100]++ |
| HASH(Str #854066) | [1335]++ [3234]-- |
| HASH(Str #854067) | [64]++ [233]-- [44]++ |
| HASH(Str #854068) | [12]++ [14]-- |
| HASH(Str #854069) | [6778]++ |



Examine

Hash(string) →

The Hypothesis

Strings have the potential for disproportionate impact on the feature vector



The Whitelist provides a hint as to what type of executables are "good" (e.g. Rocket League) and may have been used to retrain the model at a later stage



If we strip the strings from the good PEs and **carefully inject them into a malicious payload**, we may be able to fool the model, as they will overpower the effect of "negative" properties. Note that the model does not regard "attacker economics".

Note that we are NOT aiming to fool the whitelisting mechanism, rather the main model!

This would never work, right?

⋮

```
66 1B C9 66 23 C2 49 8B D7 66 81 E1 00 04 66 0B
C8 0F BF 45 0A 89 44 24 30 48 8B 45 00 66 89 4C
24 28 48 8B CE 48 89 44 24 20 E8 39 68 00 00 33
C9 48 8B D8 48 3B C1 0F 84 97 0C 00 00 0F B7 40
30 66 3B 45 0A 7D 04 66 89 45 0A 45 33 ED 41 3A
FD 74 19 0F B6 4B 3A 0F B6 84 24 89 00 00 00 41
```

```
f.Éf#ÂI<×f.á..f.
È.¿E.‰D$0H<E.f‰L
$(H<ÎH‰D$ è9h..3
ÉH<ØH;Á.,,-....·@
0f;E.}.f‰E.E3íA:
ýt..¶K:..¶,,$‰...A
```

+

```
52 75 73 73 69 61 6E 0D 0A 74 0D 0A 52 65 6D 6F
76 65 0D 0A 39 3B 75 0D 0A 44 65 6C 65 74 65 0D
0A 69 64 35 0D 0A 35 35 35 35 0D 0A 43 68 61 6E
6E 65 6C 0D 0A 64 65 73 63 72 69 78 74 69 6F 6E
0D 0A 4D 45 54 0D 0A 25 73 0D 0A 30 31 32 33 34
35 36 37 38 39 41 42 43 44 45 46 0D 0A 77 69 6E
33 32 0D 0A 23 0D 0A 25 30 32 78 0D 0A 29 0D 0A
46 54 0D 0A 25 38 64 0D 0A 31 35 33 36 0D 0A 52
45 44 0D 0A 4C 6F 67 0D 0A 31 30 31 30 0D 0A 42
```

```
Russian..t..Remo
ve..98u..Delete.
.md5..5555..Chan
nel..description
.MET..%s..01234
56789ABCDEF..win
32..#..%02x..)..
FT..%0d..1536..R
ED..Log..1010..B
```

=

Russian
t
Remove
98u
Delete
md5
5555
Channel
description
MET
%s
0123456789ABCDEF
win32

%02x
)
FT
%0d
1536
RED
Log
1010
B6
B14
UG
DLL

Let's have a look...



Summon the Malware Hordes

| Malware | Score Before | Score After |
|-----------|--------------|-------------|
| CoinMiner | -826 | 884 |
| Dridex | -999 | 996 |
| Emotet | -923 | 625 |
| Gh0stRAT | -975 | 998 |
| Kovter | -999 | 856 |
| Nanobot | 971 | 999 |
| Pushdo | -999 | 999 |
| Qakbot | -998 | 991 |
| Trickbot | -973 | 774 |
| Zeus | -997 | 997 |

Tests on 384 samples from theZoo repository:

88.54% of malware passed as benign

Average score before treatment = -0.92 (min is -1)

Average score after mutation = 0.75 (max is 1)

Average change in score = +1.67 (out of a range of 2).

Publication & Cylance's Response

July 21st, Cylance's Threat Vector

...researchers publicly disclosed a specific bypass of CylancePROTECT®. We verified the issue was not a universal bypass as reported, but rather a technique that allowed for one of the anti-malware components of the product to be bypassed in certain circumstances. The issue has been resolved for cloud-based scoring and a new agent will be rolled out to endpoints in the next few days.

We are still waiting for a fix for the SmartAV product...



Kim Zetter @KimZetter · Aug 14
Ryan Perme, founder/chief scientist at Cylance, has responded and said SmartAV (the consumer version) trails behind corporate version in update rollouts. He said the issue was fixed in the corp version of their AV and is checking now on what's going on with the consumer version.

Kim Zetter @KimZetter · Aug 14
I already reached out to Cylance and will update if I hear more.

Kim Zetter @KimZetter · Aug 14
Cylance pr told me the company fixed bypass I wrote about last month and rolled out fix to customers July 26. But researchers who discovered issue told me they installed fresh version of SmartAV agent this wk and can still bypass it using same trick



Researchers Easily Trick Cylance's AI-Based Antivirus Into Thinking ...
By taking strings from an online gaming program and appending them to malicious files, researchers were able to trick Cylance's AI-based antivirus...
vice.com

Questions?



SKYLIGHT

ZERO COMPROMISE

Thank You!

adi@skylightcyber.com