VISTA ANALYTICS

1333 H STREET, N.W. SUITE 600W
WASHINGTON, D.C. 20005
(202) 808-2286

Intrusion Detection

Background

Intrusion Detection can be defined as "...the act of detecting actions that attempt to compromise the confidentiality, integrity or availability of a resource."[1] The intrusion detector learning task is to build a predictive model (i.e. a classifier) capable of distinguishing between “bad” connections, called intrusions or attacks, and “good” normal connections. The trained model can then be deployed to flag any potential unauthorized or illicit connections.

Intrusion Detection is one part of a holistic approach to cyber security as evidenced by the wide range of threats that are currently plaguing the world’s computer systems. This analysis is not meant to ignore important components such as insider threats or sophisticated non-intrusion attacks. It is merely one part of the analysis.

Data & Features

In this article, we use machine learning to learn patterns from historical data, and to build a model that detects network intrusion. Specifically, we used a sample of NSL-KDD data set[2] which is the refined version of the KDD cup99 data set. In this data set, there are 41 features (please see feature description in Table 1) and 25192 connections. A connection is a sequence of TCP packets starting and ending at well-defined times, between which data flows to and from a source IP address to a target IP address under a well-defined protocol. Each connection is labeled as either normal, or as an attack, with exactly one specific attack type. Among these connections, 13449 (53.39%) are normal connections, 9234 (36.65%) are DoS attack connections, 2289 (9.09%) are Probe attack connections, 11 (0.04%) are U2R attack connections, and 209 (0.83%%) are R2L attack connections. Table 2 summaries the mapping between attack class and attack type.

VISTA ANALYTICS

1333 H STREET, N.W. SUITE 600W
WASHINGTON, D.C. 20005
(202) 808-2286

Table 1. Feature Description

Attribute No.	Attribute Name	Description	Sample Data
BASIC FEATURES OF EACH NETWORK CONNECTION VECTOR			
1	Duration	Length of time duration of the connection	0
2	Protocol_type	Protocol used in the connection	TCP
3	Service	Destination network service used	ftp_data
4	Flag	Status of the connection – Normal or Error	SF
5	Src_bytes	Number of data bytes transferred from source to destination in single connection	491
6	Dst_bytes	Number of data bytes transferred from destination to source in single connection	0
7	Land	if source and destination IP addresses and port numbers are equal then, this variable takes value 1 else 0	0
8	Wrong_fragment	Total number of wrong fragments in this connection	0
9	Urgent	Number of urgent packets in this connection. Urgent packets are packets with the urgent bit activated	0
CONTENT RELATED FEATURES OF EACH NETWORK CONNECTION VECTOR			
10	Hot	Number of “hot” indicators in the content such as: entering a system directory, creating programs and executing programs	0
11	Num_failed _logins	Count of failed login attempts	0
12	Logged_in	Login Status: 1 if successfully logged in; 0 otherwise	0
13	Num_compromised	Number of ``compromised'' conditions	0
14	Root_shell	1 if root shell is obtained; 0 otherwise	0
15	Su_attempted	1 if ``su root'' command attempted or used; 0 otherwise	0
16	Num_root	Number of ``root'' accesses or number of operations performed as a root in the connection	0
17	Num_file_creations	Number of file creation operations in the connection	0
18	Num_shells	Number of shell prompts	0
19	Num_access_files	Number of operations on access control files	0
20	Num_outbound_cmds	Number of outbound commands in an ftp session	0
21	Is_hot_login	1 if the login belongs to the ``hot'' list i.e., root or admin; else 0	0
22	Is_guest_login	1 if the login is a ``guest'' login; 0 otherwise	0

VISTA ANALYTICS

1333 H STREET, N.W. SUITE 600W
WASHINGTON, D.C. 20005
(202) 808-2286

TIME RELATED TRAFFIC FEATURES OF EACH NETWORK CONNECTION VECTOR			
23	Count	Number of connections to the same destination host as the current connection in the past two counts	2
24	Srv_count	Number of connections to the same service (port number) as the current connection in the past two seconds	2
25	Serror_rate	The percentage of connections that have activated the flag (4) s0, s1, s2 or s3, among the connections aggregated in count (23)	0
26	Srv_serror_rate	The percentage of connections that have activated the flag (4) s0, s1, s2 or s3, among the connections aggregated in srv_count (24)	0
27	Rerror_rate	The percentage of connections that have activated the flag (4) REJ, among the connections aggregated in count (23)	0
28	Srv_rerror_rate	The percentage of connections that have activated the flag (4) REJ, among the connections aggregated in srv_count (24)	0
29	Same_srv_rate	The percentage of connections that were to the same service, among the connections aggregated in count (23)	1
30	Diff_srv_rate	The percentage of connections that were to different services, among the connections aggregated in count (23)	0
31	Srv_diff_host_rate	The percentage of connections that were to different destination machines among the connections aggregated in srv_count (24)	0
HOST BASED TRAFFIC FEATURES IN A NETWORK CONNECTION VECTOR			
32	Dst_host_count	Number of connections having the same destination host IP address	150
33	Dst_host_srv_count	Number of connections having same port number	25
34	Dst_host_same _srv_rate	The percentage of connections that were to the same service, among the connections aggregated in dst_host_count (32)	0.17
35	Dst_host_diff_ srv_rate	The percentage of connections that were to different services, among the connections aggregated in dst_host_count (32)	0.03
36	Dst_host_same _src_port_rate	The percentage of connections that were to the same source port, among the connections aggregated in dst_host_srv_count (33)	0.17

1333 H STREET, N.W. SUITE 600W
WASHINGTON, D.C. 20005
(202) 808-2286

37	Dst_host_srv_ diff_host_rate	The percentage of connections that were to different destination machines, among the connections aggregated in dst_host_srv_count (33)	0
38	Dst_host_serror_rate	The percentage of connections that have activated the flag (4) s0, s1, s2 or s3, among the connections aggregated in dst_host_count (32)	0
39	Dst_host_srv_serror_rate	The percent of connections that have activated the flag (4) s0, s1, s2 or s3, among the connections aggregated in dst_host_srv_count (33)	0
40	Dst_host_rerror_rate	The percentage of connections that have activated the flag (4) REJ, among the connections aggregated in dst_host_count (32)	0.05
41	Dst_host_srv_rerror_rate	The percentage of connections that have activated the flag (4) REJ, among the connections aggregated in dst_host_srv_count (33)	0

Table 2. Mapping of Attack Class and Type

Attack Class	Attack Type
DoS: Denial of service	Back, Land, Neptune, Pod, Smurf, Teardrop, Apache2, Udpstorm, Processtable, Worm
Probe: Surveillance and other probing attack	Satan, Ipsweep, Nmap, Portsweep, Mscan, Saint
R2L: unauthorized access to local super user privileges	Guess_Password, Ftp_write, Imap, Phf, Multihop, Warezmaster, Warezclient, Spy, Xlock, Xsnoop, Snmpguess, Snmpgetattack, Httptunnel, Sendmail, Named
U2R: unauthorized access from a remote machine	Buffer_overflow, Loadmodule, Rootkit, Perl, Sqlattack, Xterm, Ps

VISTA ANALYTICS

1333 H STREET, N.W. SUITE 600W
WASHINGTON, D.C. 20005
(202) 808-2286

Experiment

Data scientists at Vista Analytics built **two** classifiers on this dataset using Gradient Boosting. Gradient boosting is a machine learning technique for classification and regression problems, which produces a prediction model in the form of an ensemble of weak prediction models, typically decision trees.

Binary

The first classifier focuses on differentiating "normal" and "attack" connections. Thus, it is a binary classification problem. The confusion matrix below summarizes the outcome of the first classifier in a 10-fold cross validation experiment. General speaking, our solution performed extremely well and achieved over 99.5% accuracy.

	Predict Normal	Predict Attack
True Normal	13415	34
True Attack	78	11665

Multi-Class

The second classifier focuses on differentiating "normal connection", "DoS attack", "Probe attack", "R2L attack", and "U2R attack". Thus, it is a multi-class classification problem. The confusion matrix below summarizes the outcome of our second classifier in a 10-fold cross validation experiment. General speaking, our solution achieved over 99.6% accuracy. Since U2R is very rare in this data, the recall of U2R is not ideal. However, this can be improved by the use of over-sampling in training. We will introduce how to address imbalance data issues in a separate article.

	Pred. DoS	Pred. Normal	Pred. Probe	Pred. R2L	Pred. U2R
True DoS	9229	4	1	0	0
True Normal	4	13434	8	3	0
True Probe	1	25	2261	2	0
True R2L	0	20	1	186	2
True U2R	0	6	0	2	3

VISTA ANALYTICS

1333 H STREET, N.W. SUITE 600W
WASHINGTON, D.C. 20005
(202) 808-2286

Feature Importance

It's also important to understand the importance of features in these two classifiers. Figures below demonstrate the top 10 features of multi-class and binary classifiers respectively.

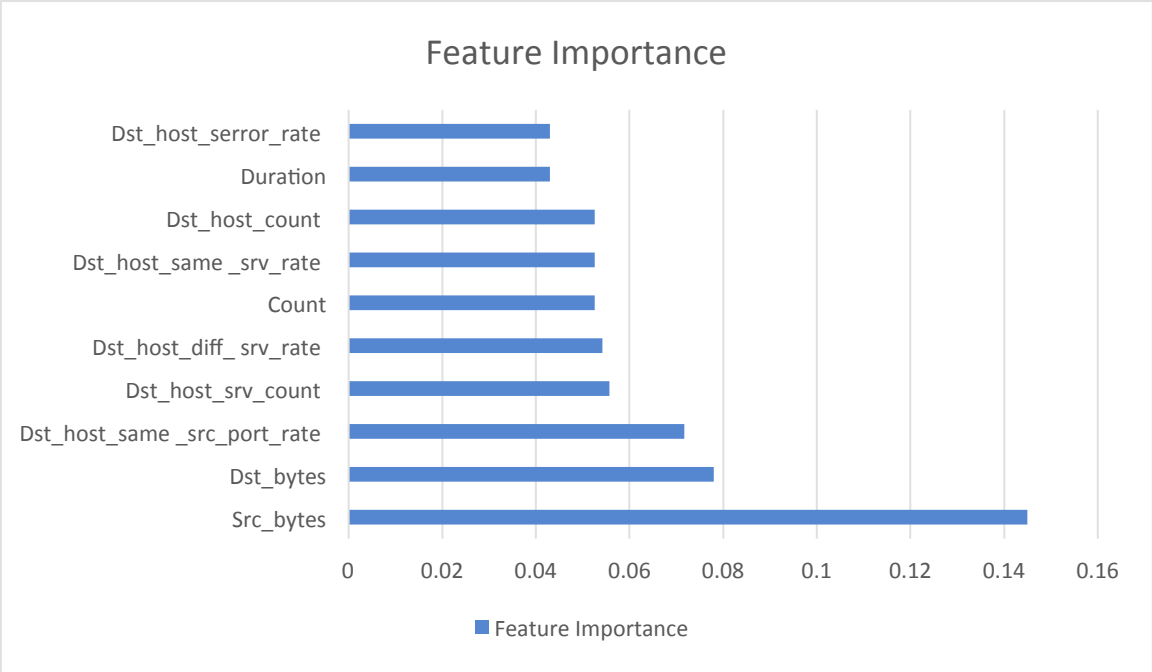

Fig. 1 Feature Importance of Multi-Class Classifier

VISTA ANALYTICS
1333 H STREET, N.W. SUITE 600W
WASHINGTON, D.C. 20005
(202) 808-2286

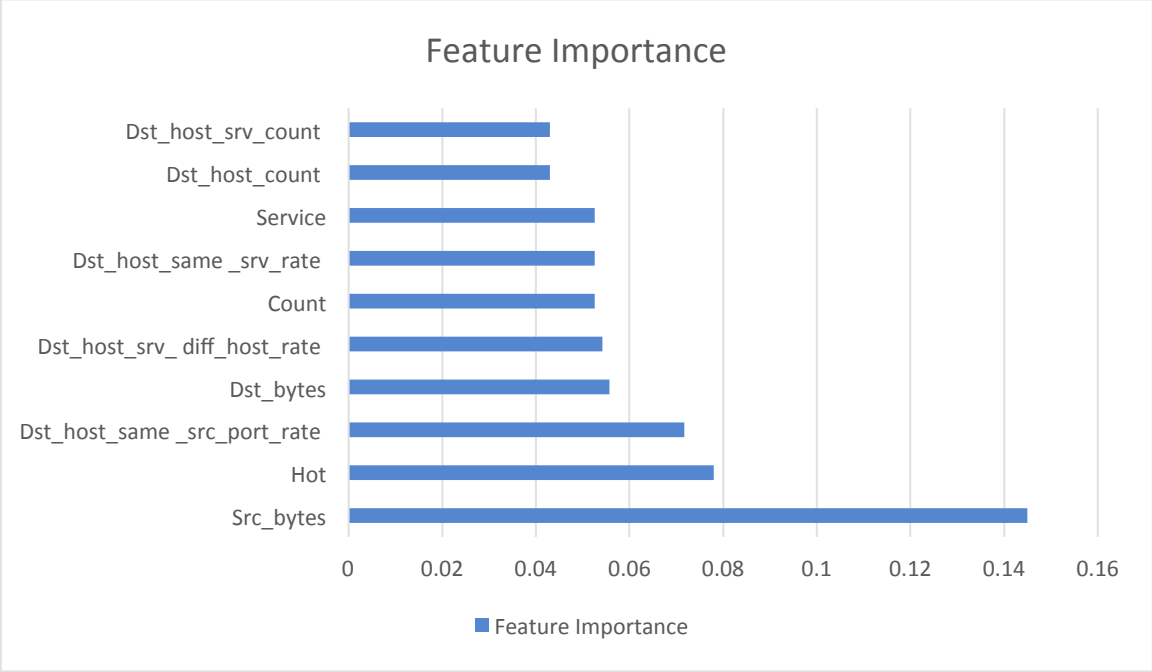

Fig. 2 Feature Importance of Binary Classifier

Reference:

[1] Wikipedia: http://en.wikipedia.org/wiki/Intrusion_detection
[2] http://www.unb.ca/research/iscx/dataset/iscx-NSL-KDD-dataset.html