

Machine Learning Implementation Security in the Wild 🌴



Denis Kolegov, Anton Nikolaev

Speakers

- Denis
 - Principal security researcher at Bi.Zone
 - Ph.D., associate professor at Tomsk State University
 - <https://twitter.com/dnkolegov>
- Anton
 - Security developer at Bi.Zone
 - Sibears CTF team player



Intro



AI Sec Project



sdnewhop.github.io/AISec/



github.com/sdnewhop/AISec/



medium.com/hackingodyssey

Contributors:

- Sergey Gordeychik
- Anton Nikolaev
- Denis Kolegov
- Maria Nedyak
- Roman Palkin



medium.com/hackingodyssey



hackingodyssey

20xx: A hacking odyssey

Region	Count
North America	13
South America	2
Europe	1983
Africa	56
Asia	16
Australia	19
Other	8, 49, 287, 189, 239, 967, 304, 6, 34

Practical Security Assessment of SD-WAN Implementations

Overview

Denis Kolegov
Oct 29 · 15 min read

AI Sec Upcoming Talks

ZeroNights 2019



Maria Nedyak (@mariya_ns)
"Hacking Medical Imaging with DICOM"

Roman Palkin (@chicken_2007)
"Malign Machine Learning Models"

Disclaimer 1/2

- This talk is by Anton and Denis
- We don't speak for our employers
- All the opinions and information here are of our responsibility

Disclaimer 2/2

This talk focuses on the implementation security aspects of ML and does not targets its specific issues such as:

1. Data poisoning attacks
2. Privacy-stealing attacks
3. Privacy-leakage attacks
4. Adversarial attacks
5. Black-box model extraction attacks
6. Physical attacks



Cloud Text-to-Speech

Text-to-speech conversion powered by machine learning.

TRY IT FREE

VIEW DOCUMENTATION

OpenCV.js Demos


- Video processing (asm.js)
- Video processing (wasm)
- Face detection (asm.js)
- Face detection (wasm)

annyang! SpeechRecognition that just works

annyang is a tiny javascript library that lets your visitors control your site with voice commands. annyang supports multiple languages, has no dependencies, weighs just 2kb and is free to use.



face-api.js playground



Speech KITT


A flexible GUI for interacting with Speech Recognition

Speech KITT makes it easy to add a GUI to sites using Speech Recognition. Whether you are using [annyang](#), a different library or `webkitSpeechRecognition` directly, KITT will take care of the GUI.

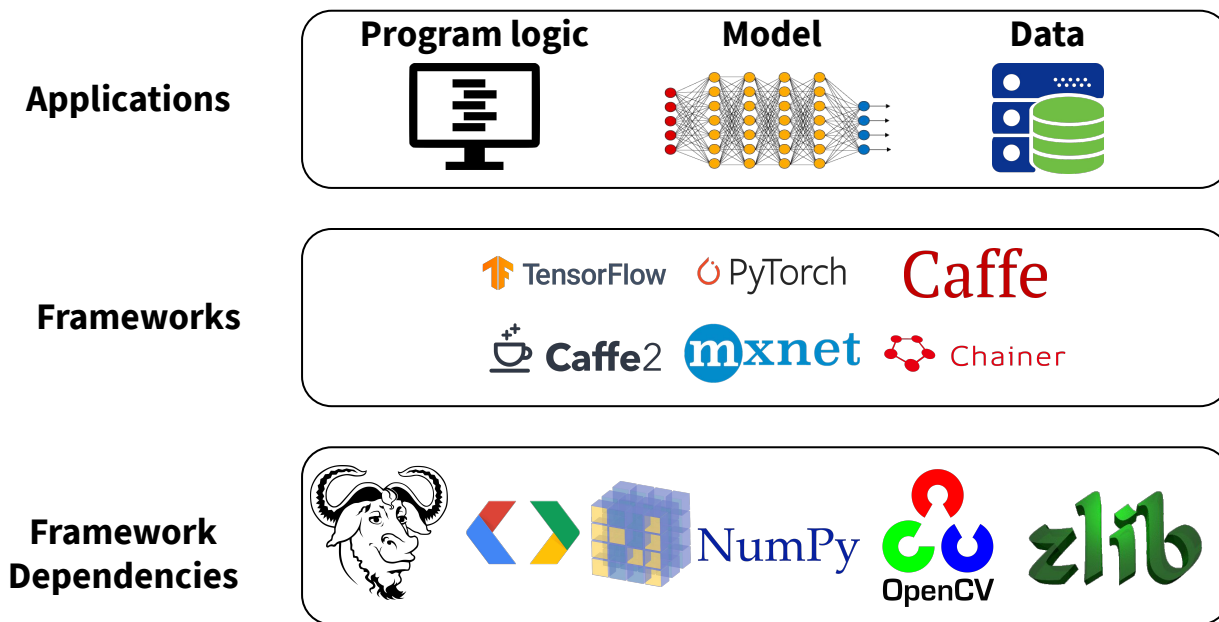
Speech KITT provides a `gui` property for the user to start or stop Speech Recognition and see its status. It also help guide the user on how to interact with your site using the `instructions` and `sampleCommands`. It can even be used to carry a `name` property for the user, asking questions the user can answer with his voice.





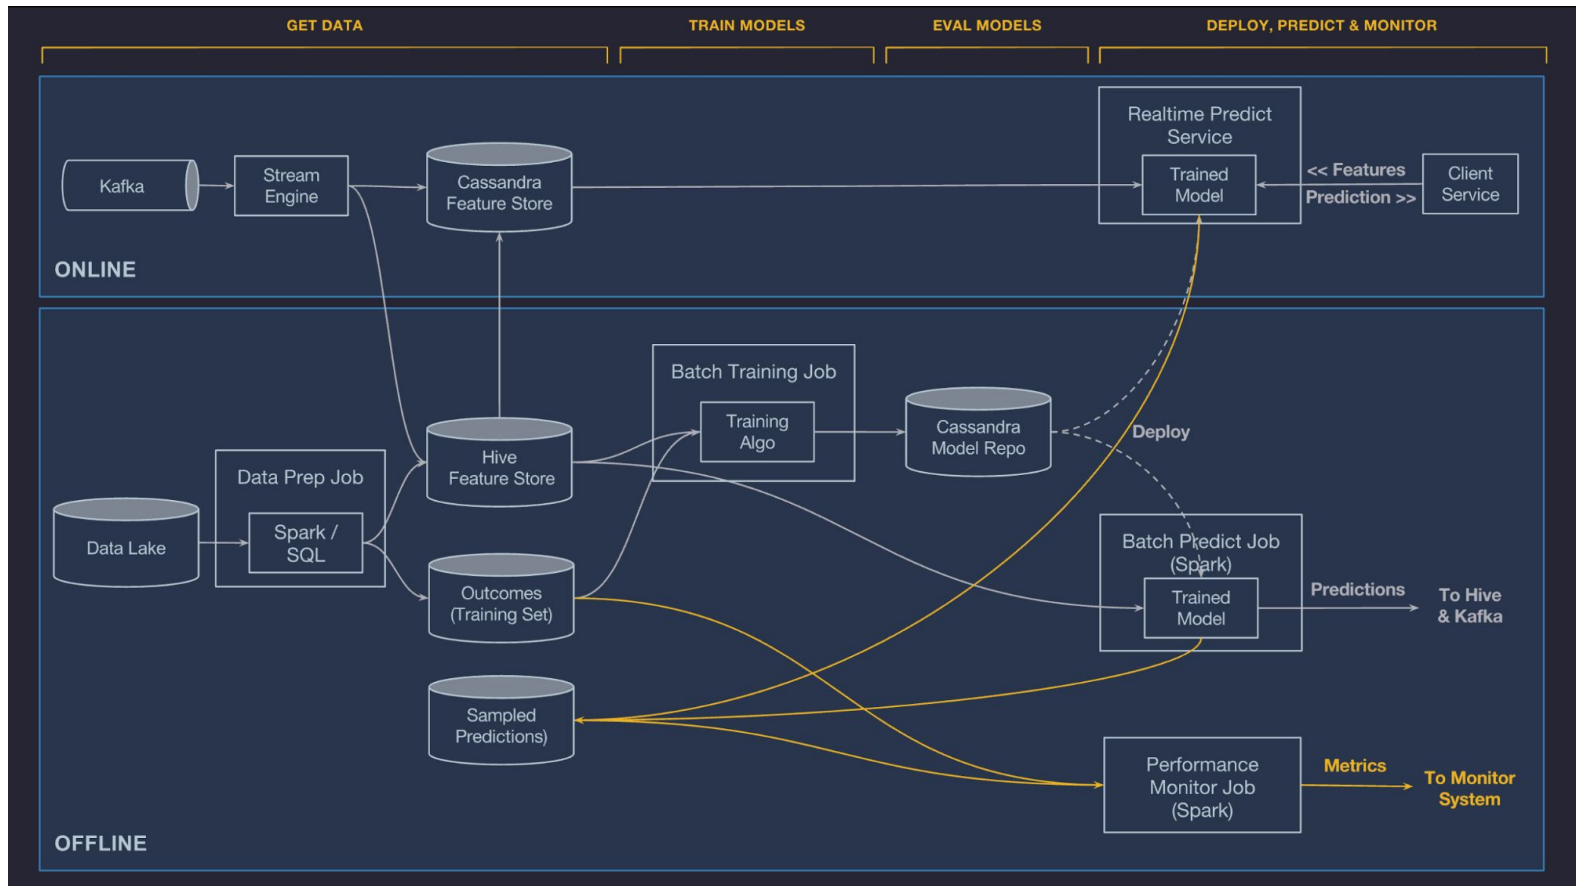
Interested In
VOICING YOUR SITE
WITH
OF CODE?



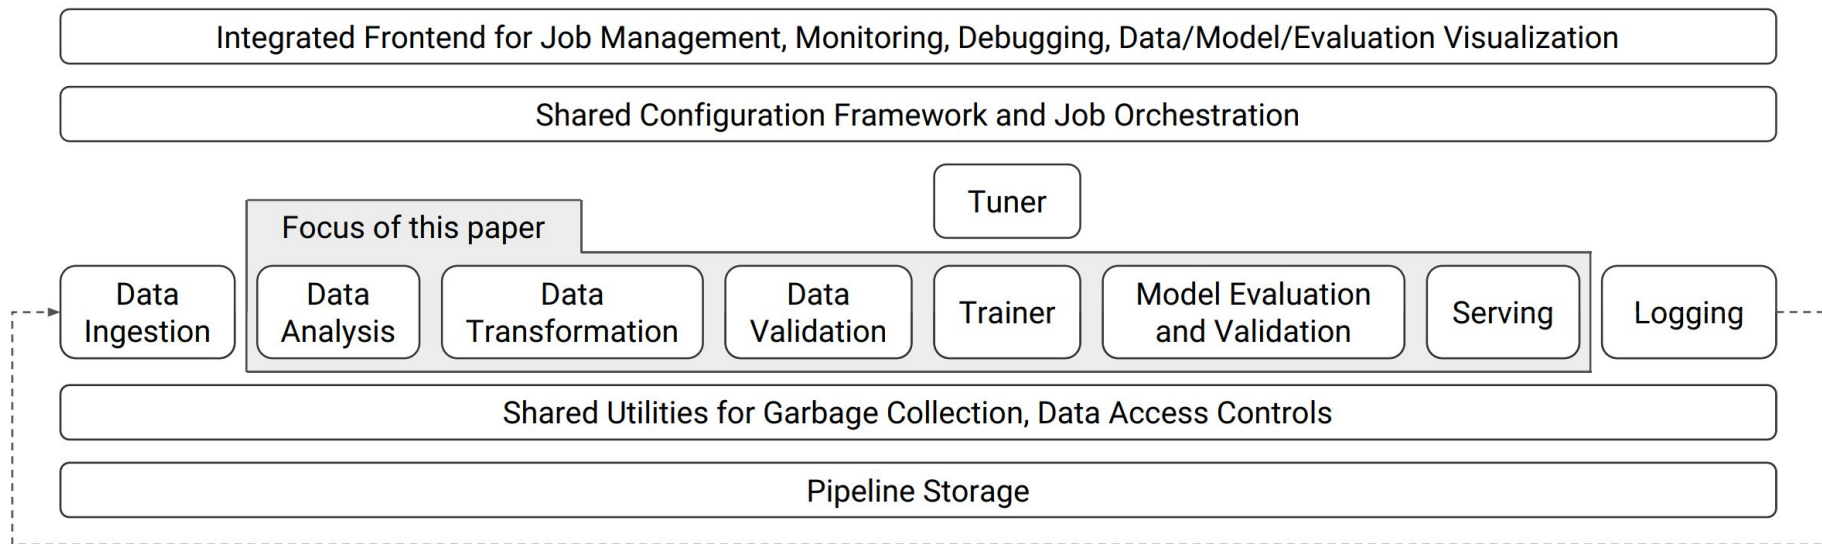
Software Layers on Machine Learning Systems

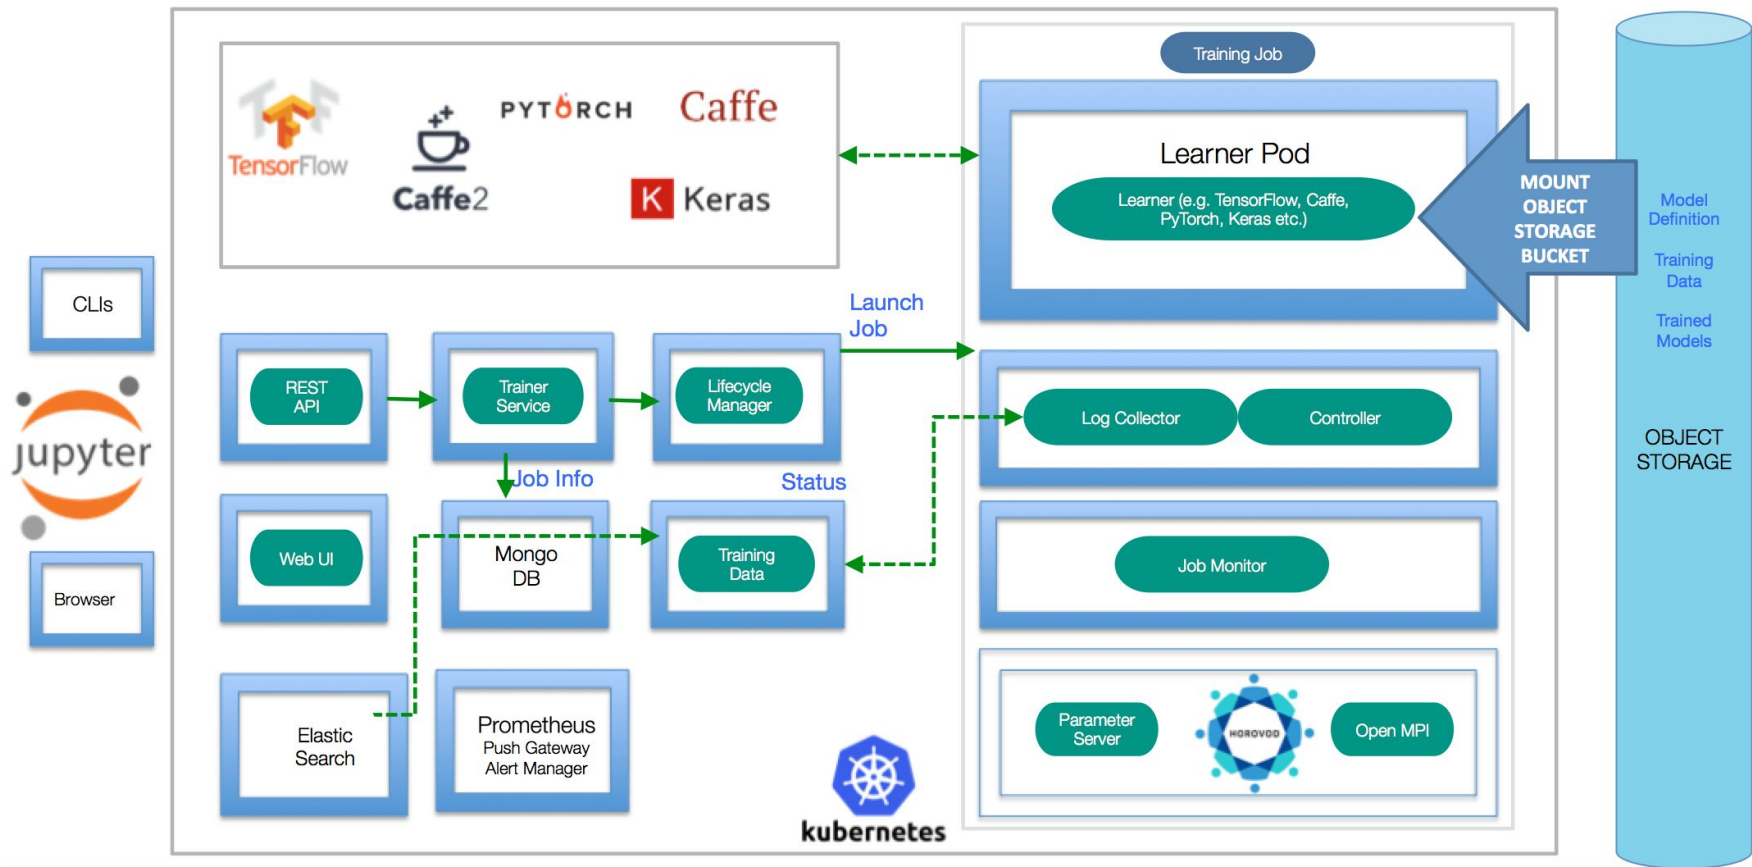






<https://github.com/1duo/awesome-ai-infrastructures>





Machine Learning Frameworks



TensorFlow



PyTorch

Caffe



Caffe2

mxnet



Chainer



Keras

theano



Cognitive
Toolkit

Frameworks Complexity and Dependencies

Framework	Files	Lines (all)	Lines of Code
Caffe	576	107.928	80.526
TensorFlow	11.219	3.310.964	2.465.296
PyTorch	5.451	1.141.599	903.697

Frame

```
name
-----
tensorflow
├── absl-py>=0.7.0
│   └── six
├── astor>=0.6.0
├── gast>=0.2.2
├── google-pasta>=0.1.6
│   └── six
├── grpcio>=1.8.6
│   └── six>=1.5.2
├── keras-applications>=1.0.8
│   ├── h5py
│   │   ├── numpy>=1.7
│   │   └── six
│   └── numpy>=1.9.1
├── keras-preprocessing>=1.0.5
│   ├── numpy>=1.9.1
│   └── six>=1.9.0
├── numpy<2.0,>=1.16.0
├── opt-einsum>=2.3.2
│   └── numpy>=1.7
├── protobuf>=3.6.1
│   ├── setuptools
│   └── six>=1.9
├── six>=1.10.0
├── tensorboard<2.1.0,>=2.0.0
│   ├── absl-py>=0.4
│   │   └── six
│   ├── google-auth-oauthlib<0.5,>=0.4.1
│   │   ├── google-auth
│   │   │   ├── cachetools<3.2,>=2.0.0
│   │   │   ├── pyasn1-modules>=0.2.1
│   │   │   │   └── pyasn1<0.5.0,>=0.4.6
│   │   │   ├── rsa<4.1,>=3.1.4
│   │   │   │   └── pyasn1>=0.1.3
│   │   │   ├── setuptools>=40.3.0
│   │   │   └── six>=1.9.0
│   │   └── requests-oauthlib>=0.7.0
│   │       ├── oauthlib>=3.0.0
│   │       └── requests>=2.0.0
│   │           ├── certifi>=2017.4.17
│   │           ├── chardet<3.1.0,>=3.0.2
│   │           ├── idna<2.9,>=2.5
│   │           └── urllib3!>=1.25.0,!<1.25.1,<1.26,>=1.21.1
│   ├── google-auth<2,>=1.6.3
│   │   ├── cachetools<3.2,>=2.0.0
│   │   ├── pyasn1-modules>=0.2.1
│   │   │   └── pyasn1<0.5.0,>=0.4.6
│   │   ├── rsa<4.1,>=3.1.4
│   │   └── pyasn1>=0.1.3
└── google-auth<2,>=1.6.3
    ├── cachetools<3.2,>=2.0.0
    ├── pyasn1-modules>=0.2.1
    │   └── pyasn1<0.5.0,>=0.4.6
    ├── rsa<4.1,>=3.1.4
    └── pyasn1>=0.1.3
```

summary

TensorFlow is an open source machine learning framework for everyone. Abseil Python Common Libraries, see <https://github.com/abseil/abseil-py>. Python 2 and 3 compatibility utilities

Read/rewrite/write Python ASTs

Python AST that abstracts the underlying Python version

pasta is an AST-based Python refactoring library

Python 2 and 3 compatibility utilities

HTTP/2-based RPC framework

Python 2 and 3 compatibility utilities

Reference implementations of popular deep learning models

Read and write HDF5 files from Python

NumPy is the fundamental package for array computing with Python. Python 2 and 3 compatibility utilities

NumPy is the fundamental package for array computing with Python. Python 2 and 3 compatibility utilities

NumPy is the fundamental package for array computing with Python. Easy data preprocessing and data augmentation for deep learning models

NumPy is the fundamental package for array computing with Python. Python 2 and 3 compatibility utilities

NumPy is the fundamental package for array computing with Python. Optimizing numpys einsum function

NumPy is the fundamental package for array computing with Python. Protocol Buffers

Easily download, build, install, upgrade, and uninstall Python packages

Python 2 and 3 compatibility utilities

Python 2 and 3 compatibility utilities

TensorBoard lets you watch Tensors Flow

Abseil Python Common Libraries, see <https://github.com/abseil/abseil-py>. Python 2 and 3 compatibility utilities

Google Authentication Library

Google Authentication Library

Extensible memoizing collections and decorators

A collection of ASN.1-based protocols modules.

ASN.1 types and codecs

Pure-Python RSA implementation

ASN.1 types and codecs

Easily download, build, install, upgrade, and uninstall Python packages

Python 2 and 3 compatibility utilities

OAuthlib authentication support for Requests.

A generic, spec-compliant, thorough implementation of the OAuth request-signing logic

Python HTTP for Humans.

Python package for providing Mozilla's CA Bundle.

Universal encoding detector for Python 2 and 3

Internationalized Domain Names in Applications (IDNA)

HTTP library with thread-safe connection pooling, file post, and more.

Google Authentication Library

Extensible memoizing collections and decorators

A collection of ASN.1-based protocols modules.

ASN.1 types and codecs

Pure-Python RSA implementation

ASN.1 types and codecs

es

$$l = \frac{\mu_1 I_1 I_2}{2\pi d} l$$

$$K = \frac{1}{2} \mu_0 \mu_r \frac{I^2}{l}$$

$$\lambda = \frac{h}{\sqrt{2eV}}$$

$$f_0 = \frac{1}{2\pi} \sqrt{\frac{g}{l}}$$

$$\oint \vec{B} \cdot d\vec{l} = \mu_0 I$$

$$C(s) = \sqrt{\frac{3kT}{m_0}} = \sqrt{\frac{3kT}{m}}$$

$$\lambda = \frac{h}{mv}$$

$$\left(\frac{E_t}{E}\right) = \frac{2 \cos^2 \theta}{1 + \cos^2 \theta}$$

$$l = \frac{\mu_1 I_1 I_2}{2\pi d} l$$

$$= \pm \sqrt{\frac{2m}{\hbar^2}} (E - V_0)$$

$$= 2\pi f$$

$$\frac{1}{\sqrt{\epsilon \cdot \mu}} = \frac{c}{\sqrt{\epsilon_r \mu_r}}$$

$$\frac{1}{v} = \frac{n_2 - n_1}{v}$$

$$\vec{S} = \vec{Q}^*$$


$$= v_e I t$$

$$= \int \frac{F_n}{R}$$


$$2^7 \cdot \dots$$

Frameworks by Language Segments


Pytorch



TensorFlow



Caffe



Common Vulnerabilities Found in TensorFlow

CVE	Vulnerability Type	Publish Date	Score	Description
CVE-2019-9635	DoS	2019-04-24	4.3	NULL pointer dereference in Google TensorFlow before 1.12.2 could cause a denial of service via an invalid GIF file.
CVE-2018-10055	Overflow	2019-04-24	5.8	Invalid memory access and/or a heap buffer overflow in the TensorFlow XLA compiler in Google TensorFlow before 1.7.1 could cause a crash or read from other parts of process memory via a crafted configuration file.
CVE-2018-8825	Overflow (Code execution)	2019-04-23	6.8	Google TensorFlow 1.7 and below is affected by: Buffer Overflow. The impact is: execute arbitrary code (local).
CVE-2018-7577	Crash	2019-04-24	5.8	Memcpy parameter overlap in Google Snappy library 1.1.4, as used in Google TensorFlow before 1.7.1, could result in a crash or read from other parts of process memory.
CVE-2018-7576	Context-dependent (Null Pointer)	2019-04-23	4.3	Google TensorFlow 1.6.x and earlier is affected by: Null Pointer Dereference. The type of exploitation is: context-dependent.
CVE-2018-7575	Overflow	2019-04-24	7.5	Google TensorFlow 1.7.x and earlier is affected by a Buffer Overflow vulnerability. The type of exploitation is context-dependent.
CVE-2018-7574	Context-dependent (Null Pointer)	2019-04-24	5.8	Google TensorFlow 1.6.x and earlier is affected by a Null Pointer Dereference vulnerability. The type of exploitation is: context-dependent.

Common Bugs Found in ML Frameworks and Dependencies

ML Framework	dep. packages	CVE-ID	Potential Threats
TensorFlow	numpy	CVE-2017-12852	DOS
TensorFlow	wave.py	CVE-2017-14144	DOS
Caffe	libjasper	CVE-2017-9782	Heap Overflow
Caffe	openEXR	CVE-2017-12596	Crash
Caffe/Torch	opencv	CVE-2017-12597	Heap Overflow
Caffe/Torch	opencv	CVE-2017-12598	Crash
Caffe/Torch	opencv	CVE-2017-12599	Crash
Caffe/Torch	opencv	CVE-2017-12600	DOS
Caffe/Torch	opencv	CVE-2017-12601	Crash

ML Framework	dep. packages	CVE-ID	Potential Threats
Caffe/Torch	opencv	CVE-2017-12602	DOS
Caffe/Torch	opencv	CVE-2017-12603	Crash
Caffe/Torch	opencv	CVE-2017-12604	Crash
Caffe/Torch	opencv	CVE-2017-12605	Crash
Caffe/Torch	opencv	CVE-2017-12606	Crash
Caffe/Torch	opencv	CVE-2017-14136	Integer Overflow

Qixue Xiao, Kang Li, Deyue Zhang, Weilin Xu,
“Security Risks in Deep Learning Implementations”
<https://arxiv.org/abs/1711.11008>

Architecture Blocks

- API Endpoints
- Training Systems
- Visualization Systems
- Infrastructure Services
- Baseboard Management Controllers
- Job and Message Queues
- Databases



TensorFlow Distributed Server + Nmap = ?

```
2019-11-03 22:14:49.809579: I tensorflow/core/distributed_runtime/rpc/grpc_channel.cc:250] Initialize
  GrpcChannelCache for job local -> {0 -> localhost:2222, 1 -> localhost:2223}
2019-11-03 22:14:49.810038: I tensorflow/core/distributed_runtime/rpc/grpc_server_lib.cc:365] Started
  server with target: grpc://localhost:2222
Starting server #0
2019-11-03 22:14:49.810084: I tensorflow/core/distributed_runtime/rpc/grpc_server_lib.cc:369] Server
  already started (target: grpc://localhost:2222)
```



TensorFlow Distributed Server + Nmap = ?

```
2019-11-03 22:14:49.809579: I tensorflow/core/distributed_runtime/rpc/grpc_channel.cc:250] Initialize
  GrpcChannelCache for job local -> {0 -> localhost:2222, 1 -> localhost:2223}
2019-11-03 22:14:49.810038: I tensorflow/core/distributed_runtime/rpc/grpc_server_lib.cc:365] Started
  server with target: grpc://localhost:2222
Starting server #0
2019-11-03 22:14:49.810084: I tensorflow/core/distributed_runtime/rpc/grpc_server_lib.cc:369] Server
  already started (target: arpc://localhost:2222)
[1] 28510 segmentation fault python3 server.py 0
(venv) → tensorflow-distributed-server
```

```
× ..ibuted-server (zsh)
→ tensorflow-distributed-server nmap localhost -p 2222
Starting Nmap 7.80 ( https://nmap.org ) at 2019-11-03 22:16 +07
Nmap scan report for localhost (127.0.0.1)
Host is up (0.00016s latency).
Other addresses for localhost (not scanned): ::1

PORT      STATE SERVICE
2222/tcp  open  EtherNetIP-1

Nmap done: 1 IP address (1 host up) scanned in 0.07 seconds
→ tensorflow-distributed-server
```



What Does it Mean?

Machine Learning applications and related infrastructure (servers, wrappers, handlers) are vulnerable to different kinds of vulnerabilities: crashes, denial of service, integer and heap overflows, etc.



What Does it Mean?

Machine Learning applications and related infrastructure (servers, wrappers, handlers) are vulnerable to different kinds of vulnerabilities: crashes, denial of service, integer and heap overflows, etc.

*And let's not forget
about control
interfaces!*



The Identified Issues

1. A huge number of the interfaces of various ML frameworks are open and accessible from the Internet
2. Most of them don't have authentication and/or access control mechanisms
3. Default credentials are not changed
4. Multiple common low-hanging fruit vulnerabilities (web, memory corruption, etc.) both on server- and client-side

The Problem

1. A huge number of people are excluded from the market
2. Most of these people are young and have low income
3. Default rates are high
4. Companies are unable to attract new clients




inaccessible

\$


d

How to Find ML Frameworks and Applications?





Search Engines

Google



Search Engi



intitle:"Kubeflow Central Dashboard"  

[All](#) [Images](#) [News](#) [Videos](#) [Maps](#) [More](#) [Settings](#) [Tools](#)

Page 3 of about 214 results (0.27 seconds)

Kubeflow Central Dashboard

██████████.205.58 ▾

Kubeflow Central Dashboard.

Notebook Servers - Kubeflow Central Dashboard

██████████.165.144 › jupyter ▾

Kubeflow Central Dashboard.

Name - Kubeflow Central Dashboard

██████████.165.144 › jupyter › new ▾

CPU / RAM. Specify the total amount of CPU and RAM reserved by your Notebook Server. For CPU-intensive workloads, you can choose more than 1 CPU (e.g. ...

Artifact Store - Kubeflow Central Dashboard

██████████.165.144 › metadata ▾

Kubeflow Central Dashboard.

Namespace memberships - Kubeflow Central Dashboard

██████████.com › jupyter › new ▾

Kubeflow Central Dashboard.

Pipelines

██████████.145.65 › pipeline ▾

... 8:45:25 AM. [Sample] Basic - Exit Handler. A pipeline that downloads a message and prints it out. Exit Handler will run at the end. For source code, refer to ...




Search Engines



How can we search
deeper?

Search Engines

FQFA



SHODAN



censys

ZoomEy 

Special Engines and Applications

"tfjs@1.0.0" Search

↓ Need more results? Try [internal pages search](#). [query syntax](#)

9 web pages in 2.56 s. URLs CSV CSV+snippets

Rank	Url	Snippets
>30M	[REDACTED]	***npm/@tensorflow/tfjs@1.0.0/dist/tf.min.js"></s
>30M	[REDACTED]	***npm/@tensorflow/tfjs@1.0.0"></script> <script
>30M	[REDACTED]	***npm/@tensorflow/tfjs@1.0.0/dist/tf.min.js"></s
>30M	[REDACTED]	***npm/@tensorflow/tfjs@1.0.0/dist/tf.min.js"></s
>30M	[REDACTED]	***npm/@tensorflow/tfjs@1.0.0/dist/tf.min.js"></s
>30M	[REDACTED]	***npm/@tensorflow/tfjs@1.0.0/dist/tf.min.js"></s
>30M	[REDACTED]	***npm/@tensorflow/tfjs@1.0.0/dist/tf.min.js"></s
>30M	[REDACTED]	***npm/@tensorflow/tfjs@1.0.0/dist/tf.min.js" typ
>30M	[REDACTED]	***npm/@tensorflow/tfjs@1.0.0/dist/tf.min.js"></s

<https://publicwww.com/>



Special Engines and Applications

searchcode

About 3,481 results: "tf.min.js"

Page 1 of 50

Filter Results

Sources

- Bitbucket 2250
- Github 988
- Google Code 411
- CodePlex 47
- GitLab 25
- Sourceforge 10
- Fedora Project 3

Languages

- Filter Languages
- Javascript 1488
 - HTML 991
 - PHP 565
 - Portable Object 200
 - Python 127
 - vim script 80
 - Lisp 56
 - Bourne Shell 22
 - diff 21
 - C# 21
 - Java 16

Makefile in ningapp-topfriends <https://github.com/ning/ningapp-topfriends.git> | 27 lines | make

```
2. CSSOPT = --type css
3. JSOPT = --type js
10. CSSSRC = css/smoothness/jquery-ui-1.7.2.custom.css css/style.css
11. CSSDEST = style.min.css
12.
13. JSSRC = tf.js
14. JSDEST = tf.min.js
15.
17.
18. ${GADGETDEST}: ${GADGETSRC} ${CSSDEST} ${JSDEST}
23.
24. ${JSDEST}: ${JSSRC}
25.     cat ${JSSRC} | ${YUIBIN} ${JSOPT} > ${JSDEST}
26.
```

app.min.js in habitrpg-mobile <https://github.com/jixing/habitrpg-mobile.git> | 61661 lines | Javascript

```
1. /*!
2. * ionic.bundle.js is a concatenation of:
3. * ionic.js, angular.js, angular-animate.js,
4. * angular-sanitize.js, angular-ui-router.js,
5. * and ionic-angular.js
6. */
15. *
16. * By @maxlynch, @benjsperry, @adambradley <3
645. // find what eventtypes we add listeners to
646. ionic.Gestures.event.determineEventTypes();
769. /**
770. * enable of disable hammer.js detection
913. * we have different events for each device/browser
914. * determine what we need and set them in the ionic.Gestures.EVE
915. */
```

<https://searchcode.com/>



Special Engines and Applications

#	Bucket	Filename	Size
1	[icon] [redacted].s3.amazonaws.com ✖	.guild/runs/2e7fbb1649f211e8bab6ee60f17f85c9/0/saved_model.pb	116.47kB
2	[icon] [redacted].s3.amazonaws.com ✖	.guild/runs/446f0cbe49b211e88d80ee60f17f85c9/0/saved_model.pb	35.37kB
3	[icon] [redacted].s3.amazonaws.com ✖	tensorflow/inception-v1/model/pipeline_tferving/0/saved_model.pb	1.63MB
4	[icon] [redacted].s3.amazonaws.com ✖	tensorflow/linear-v1/model/pipeline_tferving/1518648395/saved_model.pb	10.84kB
5	[icon] [redacted].s3.amazonaws.com ✖	tensorflow/mnist-v1/model/pipeline_tferving/0/saved_model.pb	35.37kB
6	[icon] [redacted].s3.amazonaws.com ✖	tensorflow/mnist-v2/model/pipeline_tferving/0/saved_model.pb	116.47kB
7	[icon] [redacted].s3.amazonaws.com ✖	tensorflow/mnist-v3/model/pipeline_tferving/1519517992/saved_model.pb	33.37kB
8	[icon] [redacted].s3.amazonaws.com ✖	web_model/tensorflowjs_model.pb	1.45kB
9	[icon] [redacted].test.s3.amazonaws.com ✖	web_model/tensorflowjs_model.pb	1.43kB

<https://buckets.grayhatwarfare.com/>

Main Search Problem

How to combine all possible variations of results from different search engines and special applications?

Main Search Problem(s)

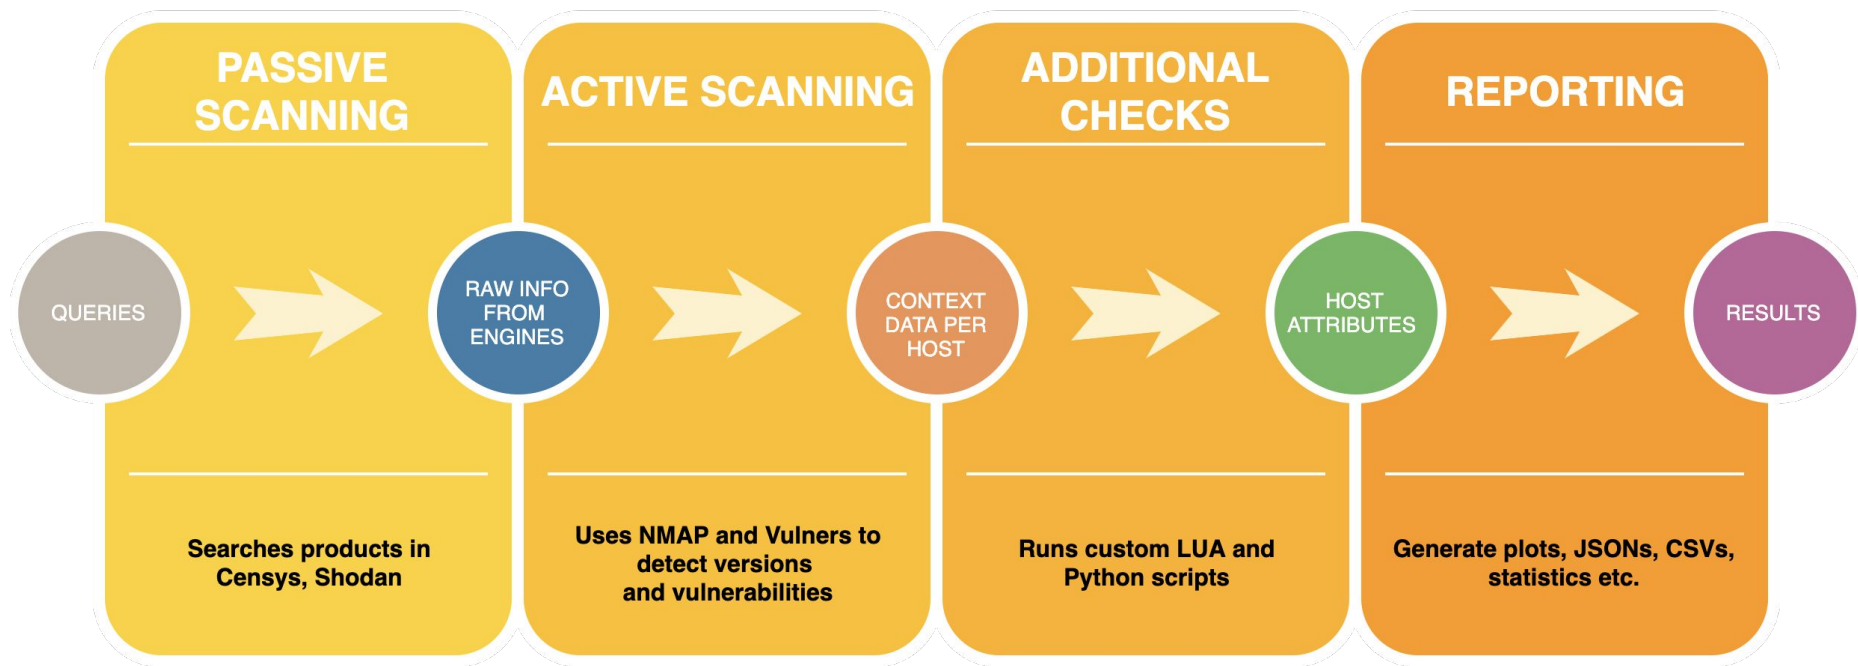
How to combine all possible variations of results from different search engines and special applications?

Moreover, how to make additional checks on them?

Main Search Problem(s)

1. How to search through different systems?
2. How to combine results from Shodan and Censys?
3. How to search for public exploits?
4. How to run custom scripts on them?
5. How to initiate ports and services scanning?
6. How to find information about vulnerabilities?
7. How to brute for SNMP public and private strings while scanning the hosts?
8. How to search hosts with some indirect methods and implicit properties?

Grinder's Workflow



Detects Vulnerabilities

```
"product": "TensorBoard",
"vendor": "Google Brain Tensorboard",
"query": "\"2016 The TensorFlow Authors\"",
"port": 8888,
"proto": "22/ssh",
"ip": [REDACTED],
"lat": 49.405,
"lng": 11.1617,
"country": "Germany",
"vulnerabilities": {
  "shodan_vulnerabilities": {},
  "vulners_vulnerabilities": {
    "CVE-2018-20852": "https://vulners.com/cve/CVE-2018-20852",
    "CVE-2019-9947": "https://vulners.com/cve/CVE-2019-9947",
    "CVE-2018-14647": "https://vulners.com/cve/CVE-2018-14647",
    "CVE-2014-4616": "https://vulners.com/cve/CVE-2014-4616",
    "CVE-2019-9636": "https://vulners.com/cve/CVE-2019-9636",
    "CVE-2019-9740": "https://vulners.com/cve/CVE-2019-9740",
    "CVE-2019-9948": "https://vulners.com/cve/CVE-2019-9948",
    "CVE-2018-1061": "https://vulners.com/cve/CVE-2018-1061",
    "CVE-2018-1060": "https://vulners.com/cve/CVE-2018-1060"
  }
},
```

```
"tcp": {
  "8888": {
    "state": "open",
    "reason": "syn-ack",
    "name": "http",
    "product": "Werkzeug httpd",
    "version": "0.12.2",
    "extrainfo": "Python 2.7.12",
    "conf": "10",
    "cpe": "cpe:/a:python:python:2.7.12",
    "script": {
      "http-title": "TensorBoard"
    }
  }
}
```

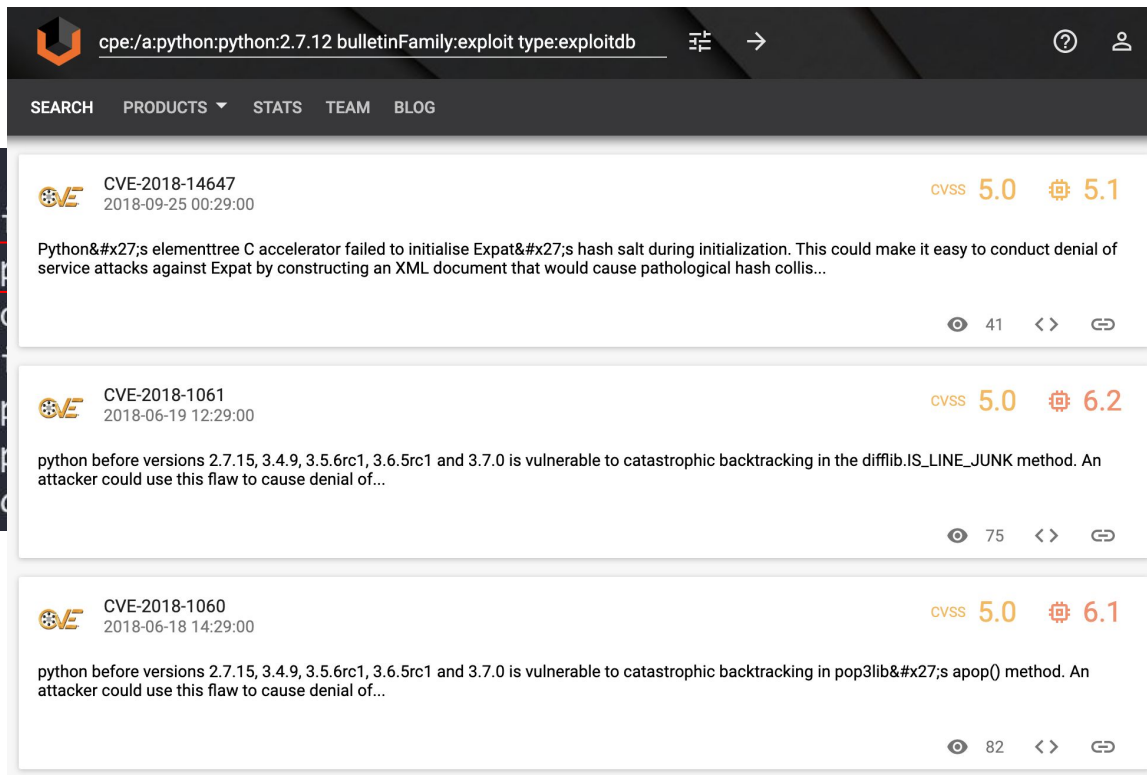
Searches Exploits

Vulners: Collect all software exploits...

- Software: cpe:/a:igor_sysoev:nginx:1.10.3, available databases: []
- Software: cpe:/a:python:python:2.7.12, available databases: ['NVD']
- Software: cpe:/a:openresty:ngx_openresty:1.15.8.1, available databases: []
- Software: cpe:/a:igor_sysoev:nginx:1.17.3, available databases: []
- Software: cpe:/a:python:python:3.7.3, available databases: ['NVD']
- Software: cpe:/a:php:php:7.3.9, available databases: []
- Software: cpe:/a:apache:http_server:2.4.25, available databases: ['NVD']




Searches Exploits

```
Vulners: Collect all
- Software: cpe:/a:python:python:2.7.12
- Software: cpe:/a:python:python:2.7.12
- Software: cpe:/a:python:python:2.7.12
- Software: cpe:/a:python:python:2.7.12
- Software: cpe:/a:python:python:2.7.12
- Software: cpe:/a:python:python:2.7.12
- Software: cpe:/a:python:python:2.7.12
```



Search results for the query: `cpe:/a:python:python:2.7.12 bulletinFamily:exploit type:exploitdb`

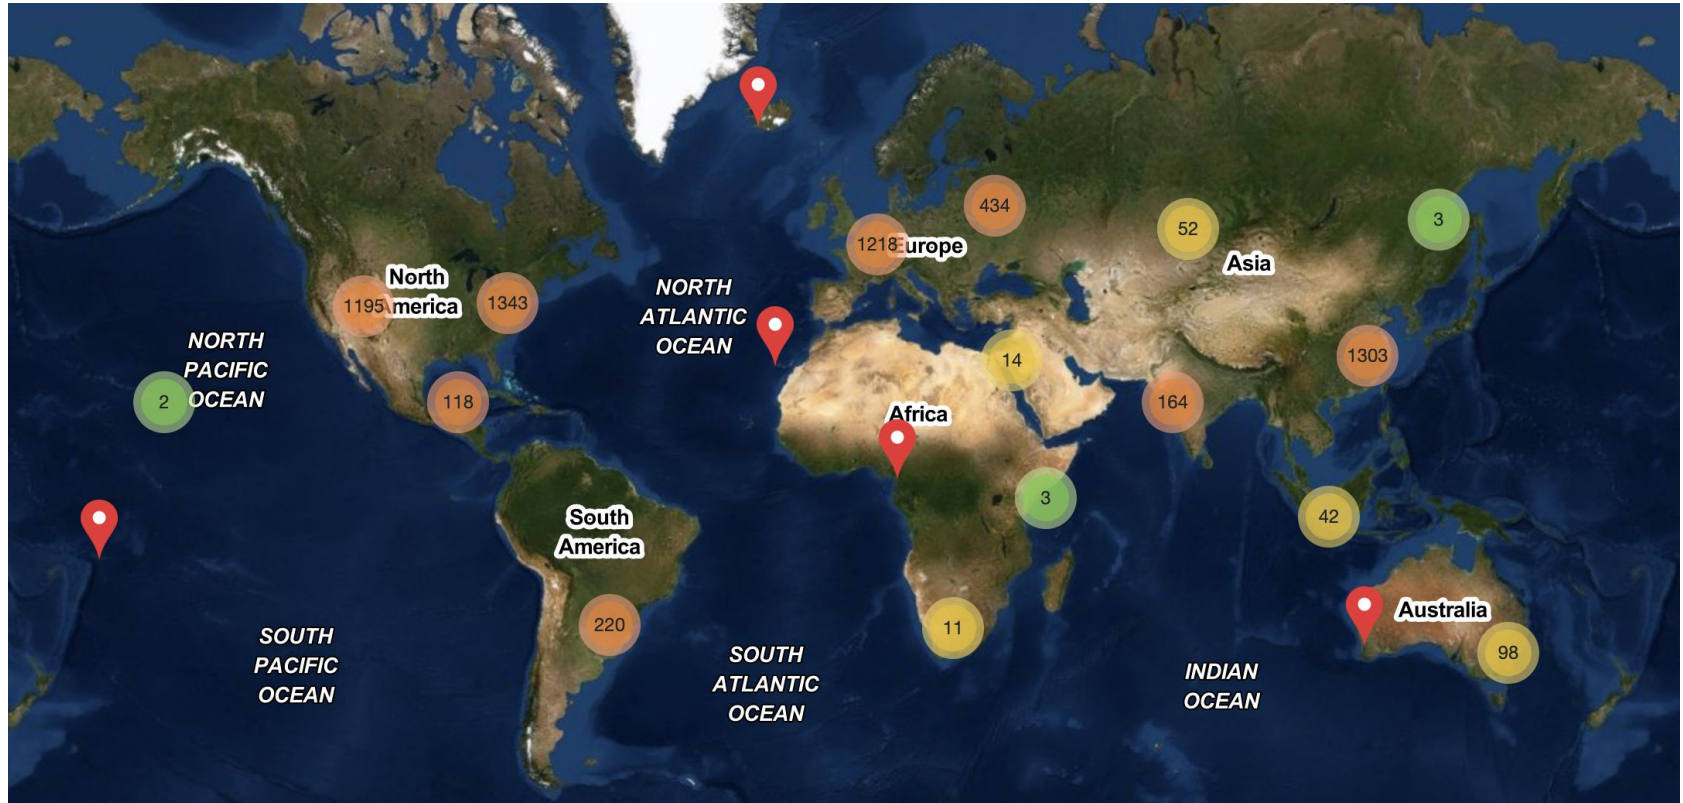
SEARCH PRODUCTS ▾ STATS TEAM BLOG

- CVE-2018-14647** (2018-09-25 00:29:00) CVSS 5.0  5.1
Python's elementtree C accelerator failed to initialise Expat's hash salt during initialization. This could make it easy to conduct denial of service attacks against Expat by constructing an XML document that would cause pathological hash collisions.
41 views
- CVE-2018-1061** (2018-06-19 12:29:00) CVSS 5.0  6.2
python before versions 2.7.15, 3.4.9, 3.5.6rc1, 3.6.5rc1 and 3.7.0 is vulnerable to catastrophic backtracking in the `difflib.IS_LINE_JUNK` method. An attacker could use this flaw to cause denial of service.
75 views
- CVE-2018-1060** (2018-06-18 14:29:00) CVSS 5.0  6.1
python before versions 2.7.15, 3.4.9, 3.5.6rc1, 3.6.5rc1 and 3.7.0 is vulnerable to catastrophic backtracking in the `pop3lib.apop()` method. An attacker could use this flaw to cause denial of service.
82 views

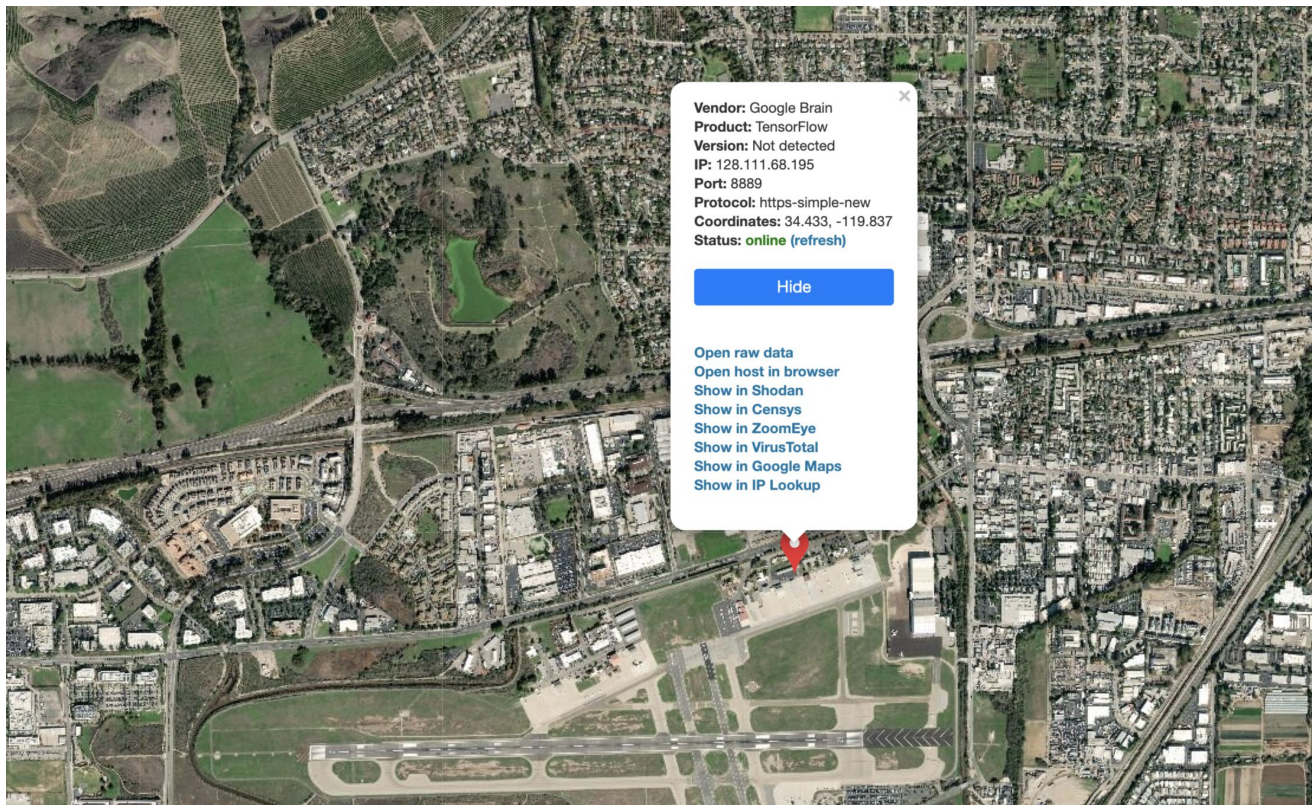
AI/ML Software Coverage

- Frameworks
 - TensorFlow
 - NVIDIA DIGITS
 - Caffe
 - TensorBoard
 - Tensorflow.js
 - brain.js
 - Predict.js
 - ml5.js
 - Keras.js
 - Figue.js
 - Natural.js
 - neataptic.js
 - ml.js
 - Clusterfck.js
 - Neuro.js
 - Deeplearn.js
 - Convnet.js
 - Synaptic.js
 - Apache mxnet
 - Databases with ML Content
 - Elasticsearch with ML data
 - MongoDB with ML data
 - Docker API with ML data
 - Databases
 - Elasticsearch
 - Kibana (Elasticsearch Visualization Plugin)
 - Gitlab
 - Samba
 - Rsync
 - Riak
 - Redis
 - Redmon (Redis Web UI)
 - Cassandra
 - Memcached
 - MongoDB
 - PostgreSQL
 - MySQL
 - Docker API
 - CouchDB
 - Job and Message Queues
 - Alibaba Group Holding AI Inference
 - Apache Kafka Consumer Offset Monitor
 - Apache Kafka Manager
 - Apache Kafka Message Broker
 - RabbitMQ Message Broker
 - Celery Distributed Task Queue
 - Gearman Job Queue Monitor
 - Interactive Voice Response (IVR)
 - ResponsiveVoice.JS
 - Inference Solutions
 - Speech Recognition
 - Speech.js
 - dictate.js
 - p5.speech.js
 - artyom.js
 - SpeechKITT
 - annyang
- ... and many more

Results



Results



Results

All the latest results can be found here:



sdnewhop.github.io/AISec/

This results including:

- Dozens of dorks for Shodan and Censys
- Interactive map (both online and offline with REST API)
- Statistics about services, ports, products, and vendors
- Statistics about vulnerabilities and exploits

Management Systems



ML Management Systems

- Training systems
 - NVIDIA DIGITS
 - MLFlow
- Visualization and tracking systems
 - TensorBoard
- Baseboard Management Controllers
 - DGX-1 Management Controller
 - DGX-2 Management Controller

Control Interfaces

mjflow

GitHub Docs

tutorial_experiment > Run 39c2c0bced8b49ad8c1917f648859a39 > train

Exploring kitti (/workspace/jobs/20190503-232021-f8f9/train_db/features) images

Show all images

Items per page: 10 - 25 - 50 - 100

0 1 2 3 4 5 ... 254



Trained Models

Select Model

Epoch #30

Download Model Make Pretrained Model Publish to inference server

Test a single image

Image Path

Upload image

Show visualizations and statistics

Classify One

Test a list of images

Upload Image List

Image folder (optix)

Number of Images

Classify Many

Top N Predictions

TensorBoard

SCALARS DEBUGGER GRAPHS DISTRIBUTIONS HISTOGRAMS

INACTIVE

Filter tags (regular expressions supported)

ModelVars

PREVIOUS PAGE NEXT PAGE

ModelVars/BoxPredictor_D/BoxPredictor/biases

ModelVars/BoxPredictor_0/BoxPredictor/weights

ModelVars/BoxPredictor_0/ClassPredictor/biases

ModelVars/BoxPredictor_0/ClassPredictor/weights

ModelVars/BoxPredictor_1/BoxPredictor/biases

ModelVars/BoxPredictor_1/BoxPredictor/weights

ModelVars/BoxPredictor_1/ClassPredictor/biases

ModelVars/BoxPredictor_1/ClassPredictor/weights



NVIDIA DIGITS

 **NVIDIA.** HIGH PERFORMANCE COMPUTING

NVIDIA DIGITS

Interactive Deep Learning GPU Training System

What's New in DIGITS 6

- Interactively train models using TensorFlow and visualize model architecture using TensorBoard
- Integrate custom plug-ins for importing special data formats such as DICOM used in medical imaging
- Pre-trained UNET model added to the DIGITS model store for image segmentation of medical images

NVIDIA DIGITS



Trained Models

Select Model

Epoch #30

Download Model Make Pretrained Model Publish to inference server

Test a single image

Image Path ?

Upload image

Browse...

Show visualizations and statistics ?

Classify One

Test a list of images

Upload Image List

Browse...

Accepts a list of filenames or urls (you can use your val.txt file)

Image folder (optional)

Relative paths in the text file will be prepended with this value before reading

Number of images use from the file

All

Leave blank to use all

Classify Many ?

Number of images to show per category

9

Top N Predictions per Category ?

NVIDIA DIGITS

Group Jobs:

Images ▾

Delete Group

Q Filter

name	framework	username	has_labels	status	elapsed	submitted
▼ Ungrouped						
VGG-16	caffe	nomoney	✘	Done	0s	Jul 13, 19
	tensorflow	nomoney	✘	Done	0s	Jul 13, 19
	caffe	nomoney	✘	Done	0s	Jul 13, 19
	tensorflow	nomoney	✘	Done	0s	Jul 13, 19
	caffe	nomoney	✘	Done	0s	Jul 13, 19
	tensorflow	nomoney	✘	Done	0s	Jul 13, 19
	caffe	nomoney	✘	Done	0s	Jul 13, 19
	caffe	nomoney	✔	Done	0s	Jul 13, 19

Job Directory

/home/nomoney/digits/digits/jobs/20190715-012323-861d

Disk Size

6.74 GB

Network

[network.py](#)

Raw tensorflow output


[tensorflow_output.log](#)

NVIDIA DIGITS

```
tensorflow_output.log x
1 WARNING: Logging before flag parsing goes to stderr.
2 W0715 01:23:25.997982 140038327420672 deprecation_wrapper.py:119] From /home/nomoney/digits/digits/tools/tensorflow/main.py:743: The name tf.app.run is deprecated. Please
  tf.compat.v1.app.run instead.
3 I0715 01:23:25.998847 140038327420672 main.py:417] Train batch size is 16 and validation batch size is 16
4 I0715 01:23:25.998908 140038327420672 main.py:421] Training epochs to be completed for each validation : 1
5 I0715 01:23:25.999357 140038327420672 main.py:425] Training epochs to be completed before taking a snapshot : 1.0
6 I0715 01:23:25.999517 140038327420672 main.py:429] Model weights will be saved as snapshot <EPOCH>_Model.ckpt
7 I0715 01:23:25.999675 140038327420672 main.py:442] Loading mean tensor from /home/nomoney/digits/digits/jobs/20190714-234820-4d43/mean.binaryproto file
8 I0715 01:23:26.007992 140038327420672 main.py:448] Loading label definitions from /home/nomoney/digits/digits/jobs/20190714-234820-4d43/labels.txt file
9 I0715 01:23:26.008080 140038327420672 main.py:454] Found 2 classes
10 I0715 01:23:26.009743 140038327420672 tf_data.py:221] Found 46 images in db /home/nomoney/digits/digits/jobs/20190714-234820-4d43/train_db
11 W0715 01:23:26.009926 140038327420672 deprecation.py:323] From /home/nomoney/digits/digits/tools/tensorflow/tf_data.py:472: string_input_producer (from
  tensorflow.python.training.input) is deprecated and will be removed in a future version.
12 Instructions for updating:
13 Queue-based input pipelines have been replaced by `tf.data`. Use `tf.data.Dataset.from_tensor_slices(string_tensor).shuffle(tf.shape(input_tensor),
  out_type=tf.int64)[0]).repeat(num_epochs)`. If `shuffle=False`, omit the `.shuffle(...)``.
14 W0715 01:23:26.013044 140038327420672 deprecation.py:323] From /home/nomoney/venv/local/lib/python2.7/site-packages/tensorflow/python/training/input.py:278: input_producer
  from tensorflow.python.training.input) is deprecated and will be removed in a future version.
15 Instructions for updating:
16 Queue-based input pipelines have been replaced by `tf.data`. Use `tf.data.Dataset.from_tensor_slices(input_tensor).shuffle(tf.shape(input_tensor), out_type=tf.int64)[0]).re
  peat(num_epochs)`. If `shuffle=False`, omit the `.shuffle(...)``.
17 W0715 01:23:26.013463 140038327420672 deprecation.py:323] From /home/nomoney/venv/local/lib/python2.7/site-packages/tensorflow/python/training/input.py:115: count_up_to (fr
  om tensorflow.python.training.input) is deprecated and will be removed in a future version.
18 Instructions for updating:
19 Queue-based input pipelines have been replaced by `tf.data`. Use `tf.data.Dataset.from_tensors(tensor).repeat(num_epochs)`. If `shuffle=False`, omit the `shuffl
  e(...)``.
20 W0715 01:23:26.015187 140038327420672 deprecation.py:323] From /home/nomoney/venv/local/lib/python2.7/site-packages/tensorflow/python/training/input.py:115: count_up_to (fr
  om tensorflow.python.ops.variables) is deprecated and will be removed in a future version.
21 Instructions for updating:
```

Logs are available

NVIDIA DIGITS



```
network.py x
1 # Preferred settings for this model is:
2 # Base Learning Rate = 0.001
3 # Crop Size = 224
4
5 from model import Tower
6 from utils import model_property
7 import tensorflow as tf
8 import tensorflow.contrib.slim as slim
9 import utils as digits
10
11 class UserModel(Tower):
12
13     @model_property
14     def inference(self):
15         x = tf.reshape(self.x, shape=[-1, self.input_shape[0], self.input_shape[1], self.input_shape[2]])
16         with slim.arg_scope([slim.conv2d, slim.fully_connected],
17                             weights_initializer=tf.contrib.layers.xavier_initializer(),
18                             weights_regularizer=slim.l2_regularizer(1e-6)):
19             model = slim.conv2d(x, 96, [11, 11], 4, padding='VALID', scope='conv1')
20             model = slim.max_pool2d(model, [3, 3], 2, scope='pool1')
21             model = slim.conv2d(model, 256, [5, 5], 1, scope='conv2')
22             model = slim.max_pool2d(model, [3, 3], 2, scope='pool2')
23             model = slim.conv2d(model, 384, [3, 3], 1, scope='conv3')
24             model = slim.conv2d(model, 384, [3, 3], 1, scope='conv4')
25             model = slim.conv2d(model, 256, [3, 3], 1, scope='conv5')
26             model = slim.max_pool2d(model, [3, 3], 2, scope='pool5')
27             model = slim.flatten(model)
28             model = slim.fully_connected(model, 4096, activation_fn=None, scope='fc1')
29             model = slim.dropout(model, 0.5, is_training=self.is_training, scope='do1')
30             model = slim.fully_connected(model, 4096, activation_fn=None, scope='fc2')
31             model = slim.dropout(model, 0.5, is_training=self.is_training, scope='do2')
32             model = slim.fully_connected(model, self.nclasses, activation_fn=None, scope='fc3')
33         return model
34
35     @model_property
36     def loss(self):
37         model = self.inference
38         loss = digits.classification_loss(model, self.y)
39         accuracy = digits.classification_accuracy(model, self.y)
40         self.summaries.append(tf.summary.scalar(accuracy.op.name, accuracy))
41         return loss
42
43
tensorflow_output.log x
1 WARNING: Logging before flag parsing goes to stderr
2 W0715 01:23:25.997982 140038327420672 deprecation_w
tf.compat.v1.app.run instead.
3 I0715 01:23:25.998847 140038327420672 main.py:417]
4 I0715 01:23:25.998908 140038327420672 main.py:421]
5 I0715 01:23:25.999357 140038327420672 main.py:425]
6 I0715 01:23:25.999517 140038327420672 main.py:429]
7 I0715 01:23:25.999675 140038327420672 main.py:442]
8 I0715 01:23:26.007992 140038327420672 main.py:448]
9 I0715 01:23:26.008080 140038327420672 main.py:454]
10 I0715 01:23:26.009743 140038327420672 main.py:462]
11 W0715 01:23:26.009926 140038327420672 main.py:466]
tensorflow.python.training.input
Instructions for updating:
Queue-based input pipelines have been replaced by a simpler
out_type=tf.int64)[0]).repeat(num_epochs) instead.
14 W0715 01:23:26.013044 140038327420672 main.py:466]
from tensorflow.python.training.input
Instructions for updating:
Queue-based input pipelines have been replaced by a simpler
(num_epochs). If 'shuffle=False', instead of
17 W0715 01:23:26.013463 140038327420672 main.py:466]
from tensorflow.python.training.input
Instructions for updating:
Queue-based input pipelines have been replaced by a simpler
19 W0715 01:23:26.015187 140038327420672 deprecation_w
tensorflow.python.ops.variables) is deprecated.
21 Instructions for updating:
```

Please

_producer

4)[0]).re

_epochs (

up to (f

Model architecture too

NVIDIA DIGITS

Exploring kitti (/workspace/jobs/20190503-232021-f8f9/train_db/features) images

[Show all images](#)

Items per page: 10 - **25** - 50 - 100

« 0 1 2 3 4 5 ... 254 »



And the train data with images



TOTAL RESULTS

55

TOP COUNTRIES



United States	13
China	12
Korea, Republic of	10
Taiwan	5
Germany	4

TOP SERVICES

Synology	31
HTTP	17
AndroMouse	3
NAS Web Interfaces	1
HTTP (8080)	1

TOP ORGANIZATIONS

Amazon.com	6
Universitaet Ulm	3
Korea Telecom	3
Taiwan Academic Network	2
Shanghai JiaoTong University	2

TOP PRODUCTS

nginx	16
-------	----

New Service: Keep track of what you have connected to the Internet. Check out [Shodan Monitor](#)

DIGITS

51.83.15.26
ip-51-83-15.eu

OVH SAS

Added on 2019-11-01 21:26:26 GMT

France

Technologies:



HTTP/1.1 200 OK

Content-Type: text/html; charset=utf-8

Content-Length: 44067

Date: Fri, 01 Nov 2019 21:26:26 GMT

DIGITS

169.44.201.108
6c.c9.2ca9.ip4.static.sl-reverse.com

SoftLayer Technologies

Added on 2019-11-03 08:07:25 GMT

United States

Technologies:



HTTP/1.1 200 OK

Content-Type: text/html; charset=utf-8

Content-Length: 44113

Date: Sun, 03 Nov 2019 08:05:04 GMT

DIGITS

175.193.50.185
Korea Telecom

Added on 2019-11-05 03:02:05 GMT

Korea, Republic of, Seoul

Technologies:



HTTP/1.1 200 OK

Server: nginx/1.10.3 (Ubuntu)

Date: Tue, 05 Nov 2019 03:02:04 GMT

Content-Type: text/html; charset=utf-8

Content-Length: 42821

Connection: keep-alive

DIGITS

202.120.39.167
Shanghai JiaoTong University

Added on 2019-11-01 22:56:33 GMT

China, Shanghai

Technologies:

HTTP/1.1 200 OK

Content-Type: text/html; charset=utf-8

Content-Length: 4378480

Date: Fri, 01 Nov 2019 22:56:32 GMT

Небезопасно [redacted] /models/20190807-113653- [refresh]

DIGITS Generic Image Model Login Info About

coco-dog-detectnet

Owner: hekun Clone Job Delete Job

Job Directory

/jobs/20190807-113653-bf61

Disk Size
160 MB

Network (train/val)
[train_val.prototxt](#)

Network (deploy)
[deploy.prototxt](#)

Network (original)
[original.prototxt](#)

Solver
[solver.prototxt](#)

Raw caffe output
[caffe_output.log](#)

Pretrained Model
/data/digits/bvlc_googlenet.caffemodel

Visualizations
[Tensorboard](#)

Dataset

coco-dog

Done Aug 07, 11:32:11 AM

- DB backend: lmdb
- Create train_db DB
 - Entry Count:** 3855
 - Feature shape:** (3, 640, 640)
 - Label shape:** (1, 52, 16)
- Create val_db DB
 - Entry Count:** 1969
 - Feature shape:** (3, 640, 640)
 - Label shape:** (1, 51, 16)

Job Status Done

- Initialized at Aug 07, 11:36:53 AM (1 second)
- Running at Aug 07, 11:36:54 AM (1 hour, 38 minutes)
- Done at Aug 07, 01:15:28 PM (Total - 1 hour, 38 minutes)

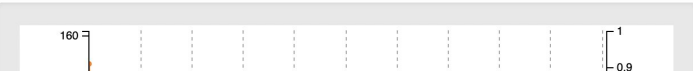
Train Caffe Model Done

Related jobs

Generic Dataset

coco-dog Done

Notes



Login Mechanism

DIGITS

[Login](#)

[Info](#) ▾

[About](#) ▾

Home

1/1 GPU available

No Jobs Running

[Datasets](#) (2)

[Models](#) (1)

[Pretrained Models](#) (0)

[New Model](#)

[Images](#) ▾

Group Jobs:

[Delete](#) [Group](#)



Filter



name	extension	framework	status	elapsed	submitted [▲]
▼ Ungrouped					
coco-dog-detectnet	image-object-detection	caffe	Done	2h	Aug 7, 19

Login Mechanism


DIGITS

[Login](#)

[Info](#) ▾

[About](#) ▾

Login

Username 

Login Mechanism

DIGITS

admin (Logout)

Info ▾

About ▾

Home

1/1 GPU available

No Jobs Running

Datasets (2)

Models (1)

Pretrained Models (0)

New Model

Images ▾

Group Jobs:

Delete Group



Filter



name	extension	framework	status	elapsed	submitted [▲]
▼ Ungrouped					
coco-dog-detectnet	image-object-detection	caffe	Done	2h	Aug 7, 19

MLFlow

Default

Experiment ID: 0

Artifact Location: ./mlruns/0

▼ Description: [🔗](#)

Search Runs:

metrics.rmse < 1 and params.model = "tree"

State:

Active ▼

Search

Filter Params:

alpha, lr

Filter Metrics:

rmse, r2

Clear

Showing 4 matching runs

Compare

Delete

Download CSV 

<input type="checkbox"/>	Date	User	Run Name	Source	Versio...	Tags	Parameters
<input type="checkbox"/>	2019-10-25 15:54:02	root		sklear...			alpha: 0.42 l1_ratio: 0.1
<input type="checkbox"/>	2019-10-24 18:22:17	root		mlflow...	7193f0		alpha: 5.0 l1_ratio: 0.1
<input type="checkbox"/>	2019-10-24 18:19:23	root		mlflow...	7193f0		alpha: 5 l1_ratio: 0.1
<input type="checkbox"/>	2019-10-24 18:06:15	root		mlflow...	7193f0		alpha: 5 l1_ratio: 0.1

MLflow is an open source platform for managing the end-to-end machine learning lifecycle. It tackles three primary functions:

- Tracking experiments to record and compare parameters and results ([MLflow Tracking](#)).
- Packaging ML code in a reusable, reproducible form in order to share with other data scientists or transfer to production ([MLflow Projects](#)).
- Managing and deploying models from a variety of ML libraries to a variety of model serving and inference platforms ([MLflow Models](#)).

MLFlow

▼ Artifacts

- ▼ models
 - MLmodel
 - conda.yaml
 - model.h5
- ▶ output
- ▼ script
 - __init__.py
 - ▶ dev
 - ▶ recipe
 - ▼ service
 - ▶ .ipynb_checkpoints
 - Dockerfile
 - build_and_push.sh
 - ▼ manage_serving
 - ▶ .ipynb_checkpoints
 - deploy_model.py
 - ▶ serve_project
 - ▶ scripts-out

Full Path: /home/ds6/submarine-project/experiment/mlflow/mlruns/1/c30e51a6a5b14af98efacfb6f62298be/artifacts/script/service/manage_serving/deploy_model.py
Size: 2.56KB

```
# environment
#pip install sagemaker
#conda install -c anaconda requests
import sagemaker as sage
from sagemaker import get_execution_role
from sagemaker.estimator import Estimator

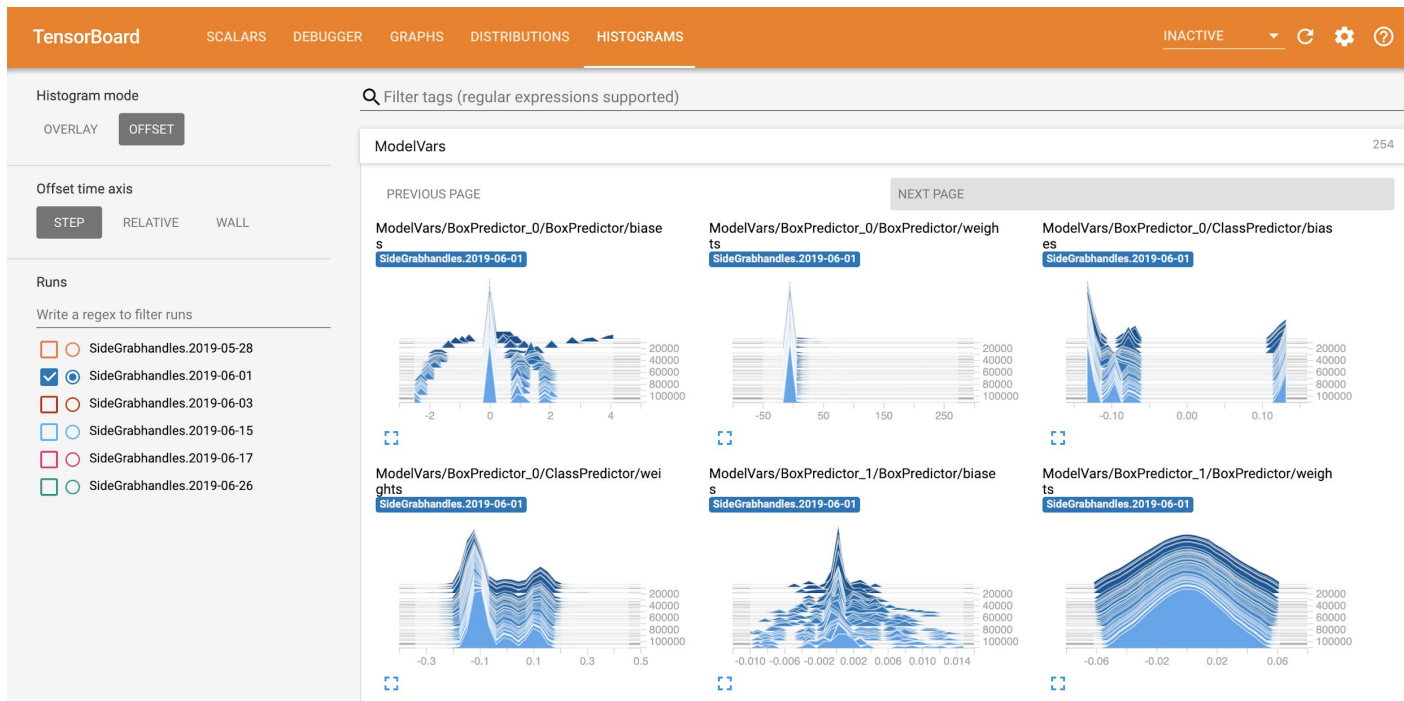
from pctcloud.data_connector import S3Connector as s3_connector
from chestxray import sagemaker_config as config

submarine_project_s3 = s3_connector(config_path=config.db_config_path, db_slug='submarine_project_s3')

#####
# Define IAM role #
#####


#####
# Create the session #
#####
algorithm_name = 'chestxray-cpu'
sess = sage.Session()
account = sess.boto_session.client('sts').get_caller_identity()['Account']
region = sess.boto_session.region_name
image = ecr_image = '{}.dkr.ecr.{}.amazonaws.com/{}:latest'.format(account, region, algorithm_name)
```

TensorBoard



TensorBoard

Branch: **master** [tensorboard](#) / [tensorboard](#) / [plugins](#) / Create new file Upload files Find file History

 **jameswex** What-If Tool progress bar and attribution sorting updates (#2892) Latest commit da9ca84 7 hours ago

..

audio	Properly handle bad requests to plugin data endpoints (#2611)	2 months ago
beholder	cleanup: use http_util.Respond (#2731)	last month
core	csp: make it globally configurable (#2756)	26 days ago
custom_scalar	Properly handle bad requests to plugin data endpoints (#2611)	2 months ago
debugger	[tensor-widget] Add colormap selection and Ctrl/Alt/Shift+wheel zoomi...	16 days ago
distribution	Promote `FrontendMetadata` from `namedtuple` to struct (#2606)	2 months ago
graph	Expose graph plugin name (#2751)	27 days ago
histogram	Place histogram bucketing logic on CPU explicitly when using TPUStrat...	2 days ago
hparams	Fix minor internal build/test failure in tf_hparams tracker (#2877)	2 days ago
image	Properly handle bad requests to plugin data endpoints (#2611)	2 months ago
interactive_inference	What-If Tool progress bar and attribution sorting updates (#2892)	7 hours ago
mesh	Revert "build: disable mesh summary v2 test (#2625)" (#2639)	23 days ago
pr_curve	core: avoid extra network request upon run selection (#2817)	13 days ago
profile	colab: use proxyPort for dynamic plugin (#2798)	15 days ago
projector	csp: fix bugs and properly treat projector (#2775)	21 days ago
scalar	core: avoid extra network request upon run selection (#2817)	13 days ago
text	Promote `FrontendMetadata` from `namedtuple` to struct (#2606)	2 months ago

TensorBoard

TensorFlow Debugger

`tfdbg` is a specialized debugger for TensorFlow. It lets you view the internal structure and states of running TensorFlow graphs during training and inference, which is difficult to debug with general-purpose debuggers such as Python's `pdb` due to TensorFlow's computation-graph paradigm.

This guide focuses on the command-line interface (CLI) of `tfdbg`. For guide on how to use the graphical user interface (GUI) of `tfdbg`, i.e., the **TensorBoard Debugger Plugin**, please visit [its README](#).

TensorBoard

TensorFlow Debug

`tfdbg` is a specialized debugger for TensorFlow graphs during training and inference, made possible due to TensorFlow's computation-graph abstraction.

This guide focuses on the command-line interface (CLI) of `tfdbg`, i.e., the **TensorBoard**

The screenshot shows the TensorBoard interface with the Debugger tab selected. The top navigation bar includes 'TensorBoard', 'SCALARS', 'DEBUGGER', 'GRAPHS', 'DISTRIBUTIONS', and 'HISTOGRAMS'. On the left, there is a 'Runtime Node List' sidebar with a 'Filter Mode' dropdown and a 'Filter Regex' input field. The main content area displays the message 'Debugger is waiting for Session.run() connections...' and three code snippets for different TensorFlow components.

```
Debugger is waiting for Session.run() connections...
```

[tf.Session:](#)

```
import tensorflow as tf
from tensorflow.python import debug as tf_debug

sess = tf.Session()
sess = tf_debug.TensorBoardDebugWrapperSession(sess, "1ef2263b814d:6000")
sess.run(my_fetches)
```

[Estimator](#) | [MonitoredSession:](#)

```
import tensorflow as tf
from tensorflow.python import debug as tf_debug

hook = tf_debug.TensorBoardDebugHook("1ef2263b814d:6000")
my_estimator.fit(x=x_data, y=y_data, steps=1000, monitors=[hook])
```


[Keras Model:](#)

```
import tensorflow as tf
from tensorflow.python import debug as tf_debug
import keras

keras.backend.set_session(
    tf_debug.TensorBoardDebugWrapperSession(tf.Session(), "1ef2263b814d:6000"))
# Define your keras model, called "model".
model.fit(...)
```

At the bottom, there is a 'Session Runs' table with columns for 'Feeds', 'Fetches', 'Targets', and '#(Devices)'. Below the table are buttons for 'STEP' and 'CONTINUE...'.

Baseboard Management Controllers



NVIDIA DGX-1 BMC

The NVIDIA® DGX-1™ Deep Learning System is the world's first purpose-built system for deep learning with fully integrated hardware and software that can be deployed quickly and easily.



1.1. Using the DGX-1: Overview

The NVIDIA DGX-1 comes with a base operating system consisting of an Ubuntu OS, Docker, Docker Engine Utility for NVIDIA GPUs, and NVIDIA drivers. This system is designed to run a number of NVIDIA-optimized deep learning framework applications packaged in Docker containers. You can use your own scheduling and management software to run jobs, and also build and run your own applications on the DGX-1.

NVIDIA DGX-1 BMC

The NVIDIA® DGX-1™ is designed for deep learning with a focus on performance, reliability, and ease of use. It is a powerful system that can be managed quickly and easily.



NVIDIA DGX-1

DU-08033-001_v23 | October 2019

User Guide

of an Ubuntu OS, users. This system is designed for network applications and management on the DGX-1.



NVIDIA DGX-1 BMC Default Credentials



Be sure to set IPMI to *Preserve* in order to preserve your BMC login credentials. If you fail to do this, the BMC username/password will be set to `qct.admin/`
`qct.admin`. If this happens, then be sure to enter the BMC dashboard and go to Configuration->Users to add a new user account and disable the qct.admin account after updating the BMC.

Also, we can try next SNMP community strings as defaults:

- qct.public
- qct.private

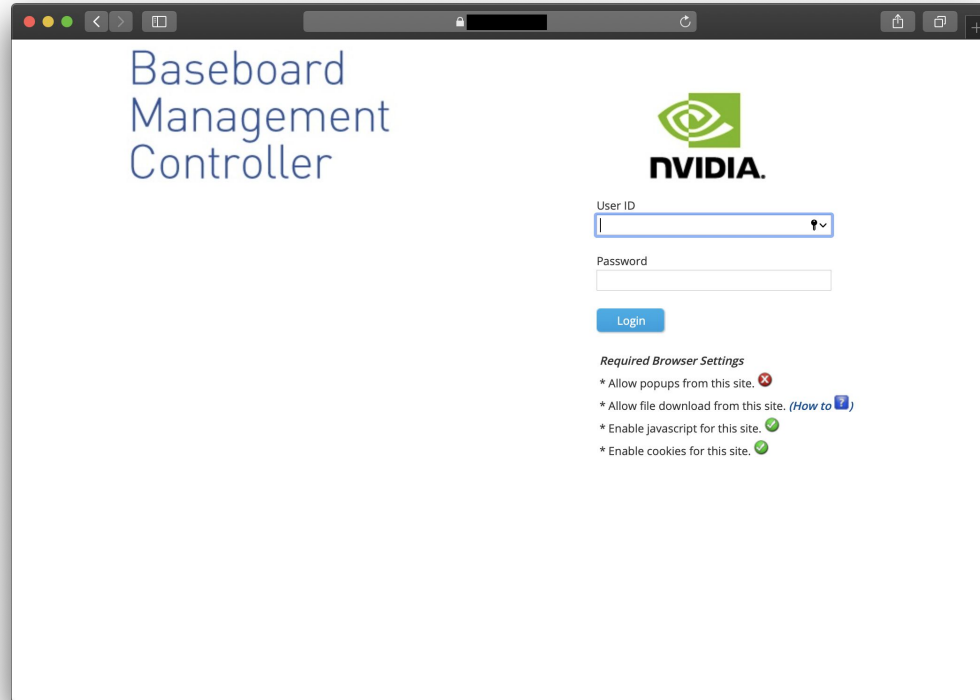
```
an.nikolaev@MacBook-Pro-Anton: ~ (zsh)
39% 38% 0.0 kB↓ 0.0 kB↑ 05.11, 7:28 PM
→ ~ snmpwalk -v 2c -c qct.private [redacted]
SNMPv2-MIB::sysDescr.0 = STRING: Linux QCTD8C4970CCA4B 3.14.17-ami #1 Sat Sep 30 14:19:55 CST 2017 armv5tejl
SNMPv2-MIB::sysObjectID.0 = OID: NET-SNMP-MIB::netSnmpAgentOIDs.10
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (293712831) 33 days, 23:52:08.31
SNMPv2-MIB::sysContact.0 = STRING: root@
SNMPv2-MIB::sysName.0 = STRING: QCTD8C4970CCA4B
SNMPv2-MIB::sysLocation.0 = STRING: Unknown
SNMPv2-MIB::sysORLastChange.0 = Timeticks: (167) 0:00:01.67
SNMPv2-MIB::sysORID.1 = OID: SNMP-FRAMEWORK-MIB::snmpFrameworkMIBCompliance
SNMPv2-MIB::sysORID.2 = OID: SNMP-MPD-MIB::snmpMPDCompliance
SNMPv2-MIB::sysORID.3 = OID: SNMP-USER-BASED-SM-MIB::usmMIBCompliance
SNMPv2-MIB::sysORID.4 = OID: SNMPv2-MIB::snmpMIB
SNMPv2-MIB::sysORID.5 = OID: TCP-MIB::tcpMIB
SNMPv2-MIB::sysORID.6 = OID: IP-MIB::ip
SNMPv2-MIB::sysORID.7 = OID: UDP-MIB::udpMIB
SNMPv2-MIB::sysORID.8 = OID: SNMP-VIEW-BASED-ACM-MIB::vacmBasicGroup
SNMPv2-MIB::sysORDescr.1 = STRING: The SNMP Management Architecture MIB.
SNMPv2-MIB::sysORDescr.2 = STRING: The MIB for Message Processing and Dispatching.
SNMPv2-MIB::sysORDescr.3 = STRING: The management information definitions for the SNMP User-based Security Model.
SNMPv2-MIB::sysORDescr.4 = STRING: The MIB module for SNMPv2 entities
SNMPv2-MIB::sysORDescr.5 = STRING: The MIB module for managing TCP implementations
SNMPv2-MIB::sysORDescr.6 = STRING: The MIB module for managing IP and ICMP implementations
SNMPv2-MIB::sysORDescr.7 = STRING: The MIB module for managing UDP implementations
SNMPv2-MIB::sysORDescr.8 = STRING: View-based Access Control Model for SNMP.
SNMPv2-MIB::sysORUpTime.1 = Timeticks: (40) 0:00:00.40
SNMPv2-MIB::sysORUpTime.2 = Timeticks: (40) 0:00:00.40
SNMPv2-MIB::sysORUpTime.3 = Timeticks: (40) 0:00:00.40
SNMPv2-MIB::sysORUpTime.4 = Timeticks: (60) 0:00:00.60
cSNMPv2-MIB::sysORUpTime.5 = Timeticks: (60) 0:00:00.60
SNMPv2-MIB::sysORUpTime.6 = Timeticks: (165) 0:00:01.65
SNMPv2-MIB::sysORUpTime.7 = Timeticks: (165) 0:00:01.65
```


BMC Default SNMP strings

```
an.nikolaev@MacBook-Pro-Anton: ~/Downloads/BMC-snmwalk-update-results/resu...
39% 39% 05.11, 7:31 PM
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.11.1 = STRING: "NVIDIA"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.11.2 = STRING: "NVIDIA"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.11.3 = STRING: "NVIDIA"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.11.4 = STRING: "NVIDIA"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.11.5 = STRING: "NVIDIA"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.12.1 = STRING: "DGX-1 with V100"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.12.2 = STRING: "DGX-1 with V100"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.12.3 = STRING: "DGX-1 with V100"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.12.4 = STRING: "DGX-1 with V100"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.12.5 = STRING: "DGX-1 with V100"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.13.1 = STRING: "1S2WU9Z0STB"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.13.2 = STRING: "N/A"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.13.3 = STRING: "N/A"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.13.4 = STRING: "N/A"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.13.5 = STRING: "N/A"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.14.1 = STRING: "v1.0"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.14.2 = STRING: "v1.0"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.14.3 = STRING: "v1.0"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.14.4 = STRING: "v1.0"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.14.5 = STRING: "v1.0"
SNMPv2-SMI::enterprises.7244.1.2.1.3.6.1.15.1 = STRING: "QTFCOU8010083"
```



NVIDIA DGX-1 BMC Interface



Open Databases with ML Data



Elasticsearch ML data



Elastic Version: 6.8.0

HTTP/1.1 200 OK
content-type: application/json; charset=UTF-8
content-length: 493



Cluster Name	elasticsearch
Status	yellow
Number of Indices	73

HTTP/1.1 200 OK
content-type: application/json; charset=UTF-8
content-length: 493

Elastic Indices:


filebeat-6.8.0-2019.08.15
ml-logs-2019-07-04
ml-logs-2019-08-23
my_index
ml-logs-2019-08-11
ml-logs-2019-08-10
ml-logs-2019-08-13
.kibana_1
ml-logs-2019-07-...

Elastic Indices:

filebeat-6.8.0-2019.08.15
ml-logs-2019-07-04
ml-logs-2019-08-23
my_index
ml-logs-2019-08-11
ml-logs-2019-08-10
ml-logs-2019-08-13
.kibana_1
ml-logs-2019-07-29
ml-logs-2019-08-14
ml-logs-2019-07-06
ml-logs-2019-07-24
ml-logs-2019-07-21
ml-logs-2019-07-20
ml-logs-2019-07-23
ml-logs-2019-07-22
ml-logs-2019-07-10
metricbeat-6.8.0-2019.08.22
metricbeat-6.8.0-2019.08.21
ml-logs-2019-06-29
ml-logs-2019-08-12
ml-logs-2019-07-15
ml-logs-2019-07-16
ml-logs-2019-07-17
ml-logs-2019-07-18
ml-logs-2019-07-19

```
{
  "name": "_KWPAUi",
  "cluster_name": "elasticsearch",
  "cluster_uuid": "-7tyf-mCSNC",
  "version": {
    "number": "6.8.0",
    "build_flavor": "default",
    "build_type": "deb",
    "build_hash": "65b6179",
    "build_date": "2019-05-15T20:06:13.172855Z",
    "build_snapshot": false,
    "lucene_version": "7.7.0",
    "minimum_wire_compatibility_version": "5.6.0",
    "minimum_index_compatibility_version": "5.0.0"
  },
  "tagline": "You Know, for Search"
}
```

MongoDB Datasets Data



Database Name	Size
datasets	28.0 GB
admin	112.0 kB
local	84.0 kB
config	60.0 kB

MongoDB Server Information

```
{
  "metrics": {
    "commands": {
      "updateUser": {
        "failed": 0,
        "total": 0
      },
      "killAllSessions": {
        "failed": 0,
        "total": 0
      },
      "dropRole"...
```

```
{
  "totalSize": 30076915712.0,
  "ok": 1.0,
  "databases": [
    {
      "sizeOnDisk": 114688.0,
      "collections": [],
      "name": "admin",
      "empty": false
    },
    {
      "sizeOnDisk": 61440.0,
      "collections": [],
      "name": "config",
      "empty": false
    },
    {
      "sizeOnDisk": 30076653568.0,
      "collections": [],
      "name": "datasets",
      "empty": false
    },
    {
      "sizeOnDisk": 86016.0,
      "collections": [],
      "name": "local",
      "empty": false
    }
  ]
},
```

MongoDB Datasets Data

```
> show dbs
admin      0.000GB
config    0.000GB
datasets  29.360GB
local     0.000GB
> use datasets
switched to db datasets
> show collections
fs.chunks
fs.files
images
scenes
test
> db.scenes.find().limit(5);
{ "_id" : ObjectId("5ca076463c1864186221d843"), "geo" : { "country" : "Russia", "region" : null, "city" : "Moscow" }, "gs" : "gs" }
{ "_id" : ObjectId("5ca076463c1864186221d844"), "geo" : { "country" : "Russia", "region" : null, "city" : "Moscow" }, "gs" : "gs" }
{ "_id" : ObjectId("5ca076463c1864186221d845"), "geo" : { "country" : "Belgium", "region" : null, "city" : "Ghent" }, "gs" : "gs" }
{ "_id" : ObjectId("5ca076463c1864186221d846"), "geo" : { "country" : "Czech_Republic", "region" : null, "city" : "Prague" }, "gs" : "gs" }
{ "_id" : ObjectId("5ca076463c1864186221d847"), "geo" : { "country" : "Czech_Republic", "region" : null, "city" : "Prague" }, "gs" : "gs" }
> db.images.find().limit(1);
{ "_id" : ObjectId("5ca220b63c18643d15a6d979"), "image_id" : ObjectId("5ca220b53c18643d15a6d977"), "image_name" : "139.3332", "xmin" : 115.50024000000002, "xmax" : 178.50024, "ymin" : 105.3336, "ymax" : 165.3336 }, { "xmin" : 613.1246499, "xmax" : 496.125, "ymin" : 333.33311999999995, "ymax" : 479.99999999999994 }, { "xmin" : 23.625360000000000 : 139.3332, "ymax" : 239.33304 }, { "xmin" : 52.874999999999999, "xmax" : 340.31268, "ymin" : 218.6668800000000 " : 437.33328 }, { "xmin" : 277.31304, "xmax" : 345.93768, "ymin" : 111.99983999999999, "ymax" : 173.33327999 : { "xmin" : 385.87536, "xmax" : 514.68768, "ymin" : 144.66672, "ymax" : 240.66672 }, { "xmin" : 458.43768, "x : 303.75, "ymin" : 93.33312000000001, "ymax" : 141.33311999999998 } ], "metadata" : { }, "scene_id" : 0
>
```

```
> show dbs
admin      0.000GB
config    0.000GB
datasets  29.360GB
local     0.000GB
> use datasets
switched to db datasets
> show collections
fs.chunks
fs.files
images
scenes
test
```

Running Containers



Running Docker Containers with ML Frameworks



Docker Version: 18.09.2

HTTP/1.1 404 Not Found
Content-Type: application/json
Date: Sun, 01 Sep 2019 21:10:17 GMT
Content-Length: 29

Docker Containers:

Image: mxschen/ai-proxy:latest
Command: /ai-serving/bin/proxy

Image: auto_pilot_w_proxy:c5
Command: /container/container_entry.sh **pytorch-container** /container/server.py

Image: mxschen/ai-proxy:latest
Command: /ai-serving/bin/proxy

Image: auto_pilot_w_proxy:c3
Command: /container/container_entry.sh **tensorflow-container** /container/server.py

Image: mxschen/ai-proxy:latest
Command: /ai-serving/bin/proxy


Image: mxschen/ai-proxy:latest
Command: /ai-serving/bin/proxy

Image: auto_pilot_w_proxy:c1
Command: /container/container_entry.sh **tensorflow-container** /container/server.py

Running Docker Containers with ML Frameworks

- **GET** `/containers/json`
to list all of the containers
- **GET** `/images/json`
to list all of the images
- **GET** `/containers/(id or name)/logs`
to export logs
- **GET** `/containers/(id or name)/export`
to fully export some container by id

```
{
  {
    "Id": "6486172d69aec4bfc49d2a1e29925cc2cc0ef7513df8aa3561b620c213d5e0b5",
    "Names": [
      "/tensorrt"
    ],
    "Image": "nvcr.io/nvidia/tensorflow:19.07-py3",
    "ImageID": "sha256:a9c6af3f3056e8bbc13dee8d676bba6f55a750f36abd76d6944b0fc36ad1f709",
    "Command": "/usr/local/bin/nvidia_entrypoint.sh bash",
    "Created": 1571408583,
    "Ports": [
      {
        "IP": "0.0.0.0",
        "PrivatePort": 6006,
        "PublicPort": 6006,
        "Type": "tcp"
      },
      {
        "PrivatePort": 6064,
        "Type": "tcp"
      },
      {
        "IP": "0.0.0.0",
        "PrivatePort": 8888,
        "PublicPort": 8888,
        "Type": "tcp"
      }
    ]
  }
},
]
```



Medical Imaging



NVIDIA AI Annotation Assistance API


NVIDIA AIAA Server documentation

REDACTED

gps.topo.auth.gr

Aristotle University of Thessaloniki

Added on 2019-10-02 01:47:28 GMT

 Greece, Thessaloniki

Technologies: 

HTTP/1.0 200 OK

Content-Length: 153094

Content-Type: text/html; charset=utf-8

Last-Modified: Mon, 19 Aug 2019 20:57:10 GMT

Cache-Control: max-age=43200, public

Expires: Wed, 02 Oct 2019 13:47:27 GMT

ETag: "1566248230.0-153094-234361760"

Server: Werkzeug/0.15.5 Python/3.5.2

Date: Wed, 02 0...

NVIDIA AI Annotation Assistance API

NVIDIA AIAA Server documentation

REDACTED

gps.topo.auth.gr

Aristotle University of Thessaloniki

Added on 2019-10-

 Greece, The

Technologies: 

HTTP/1.0 200 OK

Content-Length: 153094

AI Annotation Assistance server API 1.0.0 OAS3

</docs/openapi.yaml>

NVIDIA Deep Learning for Medical Imaging. Artificial Intelligence Annotation Assistance server API specification. This specification defines inference and smart polygon API. [Try/Visualize APIs](#)

API (v1)

GET `/v1/models` Retrieve the list of available models

POST `/v1/dextr3d` Request Annotation on 3D NIFTI image

POST `/v1/segmentation` Request Segmentation on 3D NIFTI image

POST `/v1/mask2polygon` Convert a 3D mask into slices of 2D polygons

POST `/v1/fixpolygon` Adjust polygons to a better-fit 2D/3D polygons

NVIDIA AI AA

```
▼ [
  ▼ {
    ▼ "roi": [
      128,
      128,
      128
    ],
    "name": "annotation_ct_liver",
    "sigma": 3,
    "version": "1",
    ▼ "labels": [
      "liver"
    ],
    "description": "A pre-trained model for volumetric (3D) annotation of the liver in portal venous phase CT image",
    "type": "annotation",
    "internal name": "annotation_ct_liver",
    "padding": 20
  },
]
```

What is the problem here?

In some cases, medical ML frameworks and AI systems are connected with PACS servers, which is a medical imaging and archiving technology.

How to retrieve information from PACS?

Q.4 DIMSE-C C-FIND Service

[Prev](#)**Q Relevant Patient Information Query Service Class (Normative)**[Next](#)

Q.4 DIMSE-C C-FIND Service

The DIMSE-C C-FIND service is the operation by which relevant patient information is queried and provided.

How to retrieve information from PACS?

Q.4 DIMSE-C C-FIND Service		
Prev	Q Relevant Patient Information Query Service Class (Normative)	Next

Q.4 DIMSE-C C-FIND Service

The DIMSE-C C-FIND service is the operation by which relevant patient information is queried and provided.

C.2.2.2.4 Wild Card Matching		
Prev	C.2.2.2 Attribute Matching	Next

C.2.2.2.4 Wild Card Matching

If the Attribute is not a date, time, signed long, signed short, unsigned short, unsigned long, floating point single, floating point double, other byte string, other word string, unknown, Attribute tag, decimal string, integer string, age string or UID and the value specified in the request contains any occurrence of an "*" or a "?", then "*" shall match any sequence of characters (including a zero length value) and "?" shall match any single character. This matching is case sensitive, except for Attributes with a PN Value Representation (e.g., Patient Name (0010,0010)).

For Attributes with a PN value representation, including the case of extended negotiation of fuzzy semantic matching, wild card matching is implementation dependent and shall be specified in the conformance statement.

How to retrieve information from PACS?

Q.4 DIMSE-C C-FIND Service		
Prev	Q Relevant Patient Information Query Service Class (Normative)	Next

Q.4 DIMSE-C C-FIND Service

The DIMSE-C C-FIND service is the operation by which relevant patient information is queried and provided.

C.2.2.2.4 Wild Card Matching		
Prev	C.2.2.2 Attribute Matching	Next

C.2.2.2.4 Wild Card Matching

If the Attribute is not a date, time, signed long, signed short, unsigned short, unsigned long, floating point single, floating point double, other byte string, other word string, unknown, Attribute tag, decimal string, integer string, age string or UID and the value specified in the request contains a "*" or a "?", then "*" shall match any sequence of characters (including a zero length value) and "?" shall match any single character. For Attributes with a PN Value Representation (e.g., Patient Name (0010,0010)).

For Attributes with a PN value representation, including the case of extended negotiation of fuzzy matching, the "*" shall match any sequence of characters and shall be specified in the conformance statement.



```
C-FIND query status: 0xff00
(0008, 0000) Group Length UL: 44
(0008, 0052) Query/Retrieve Level CS: 'PATIENT'
(0008, 0054) Retrieve AE Title AE: 'dicom\x00'
(0008, 0056) Instance Availability CS: 'ONLINE'
(0010, 0000) Group Length UL: 60
(0010, 0010) Patient's Name PN: 'ARCHIPOVA G.V.'
(0010, 0020) Patient ID LO: '295'
(0010, 0030) Patient's Birth Date DA: '19370113'
(0010, 0040) Patient's Sex CS: 'F'
C-FIND query status: 0xff00
(0008, 0000) Group Length UL: 44
(0008, 0052) Query/Retrieve Level CS: 'PATIENT'
(0008, 0054) Retrieve AE Title AE: 'dicom\x00'
(0008, 0056) Instance Availability CS: 'ONLINE'
(0010, 0000) Group Length UL: 76
(0010, 0010) Patient's Name PN: 'ARIKAINEN V.A.'
(0010, 0020) Patient ID LO: 'K26032.ARIKA.194701'
(0010, 0030) Patient's Birth Date DA: '19470119'
(0010, 0040) Patient's Sex CS: 'M'
C-FIND query status: 0xff00
(0008, 0000) Group Length UL: 44
(0008, 0052) Query/Retrieve Level CS: 'PATIENT'
(0008, 0054) Retrieve AE Title AE: 'dicom\x00'
(0008, 0056) Instance Availability CS: 'ONLINE'
(0010, 0000) Group Length UL: 74
(0010, 0010) Patient's Name PN: 'ARKHIPENKO S.I.'
(0010, 0020) Patient ID LO: '1051820826000140'
(0010, 0030) Patient's Birth Date DA: '19710823'
(0010, 0040) Patient's Sex CS: 'M'
C-FIND query status: 0xff00
(0008, 0000) Group Length UL: 44
(0008, 0052) Query/Retrieve Level CS: 'PATIENT'
(0008, 0054) Retrieve AE Title AE: 'dicom\x00'
(0008, 0056) Instance Availability CS: 'ONLINE'
(0010, 0000) Group Length UL: 76
(0010, 0010) Patient's Name PN: 'ARKHIPENKO T.F.'
(0010, 0020) Patient ID LO: '1118.ARKHI.193504'
(0010, 0030) Patient's Birth Date DA: '19350425'
(0010, 0040) Patient's Sex CS: 'F'
C-FIND query status: 0xff00
(0008, 0000) Group Length UL: 44
(0008, 0052) Query/Retrieve Level CS: 'PATIENT'
(0008, 0054) Retrieve AE Title AE: 'dicom\x00'
(0008, 0056) Instance Availability CS: 'ONLINE'
(0010, 0000) Group Length UL: 74
```

We will receive a complete list of all patients

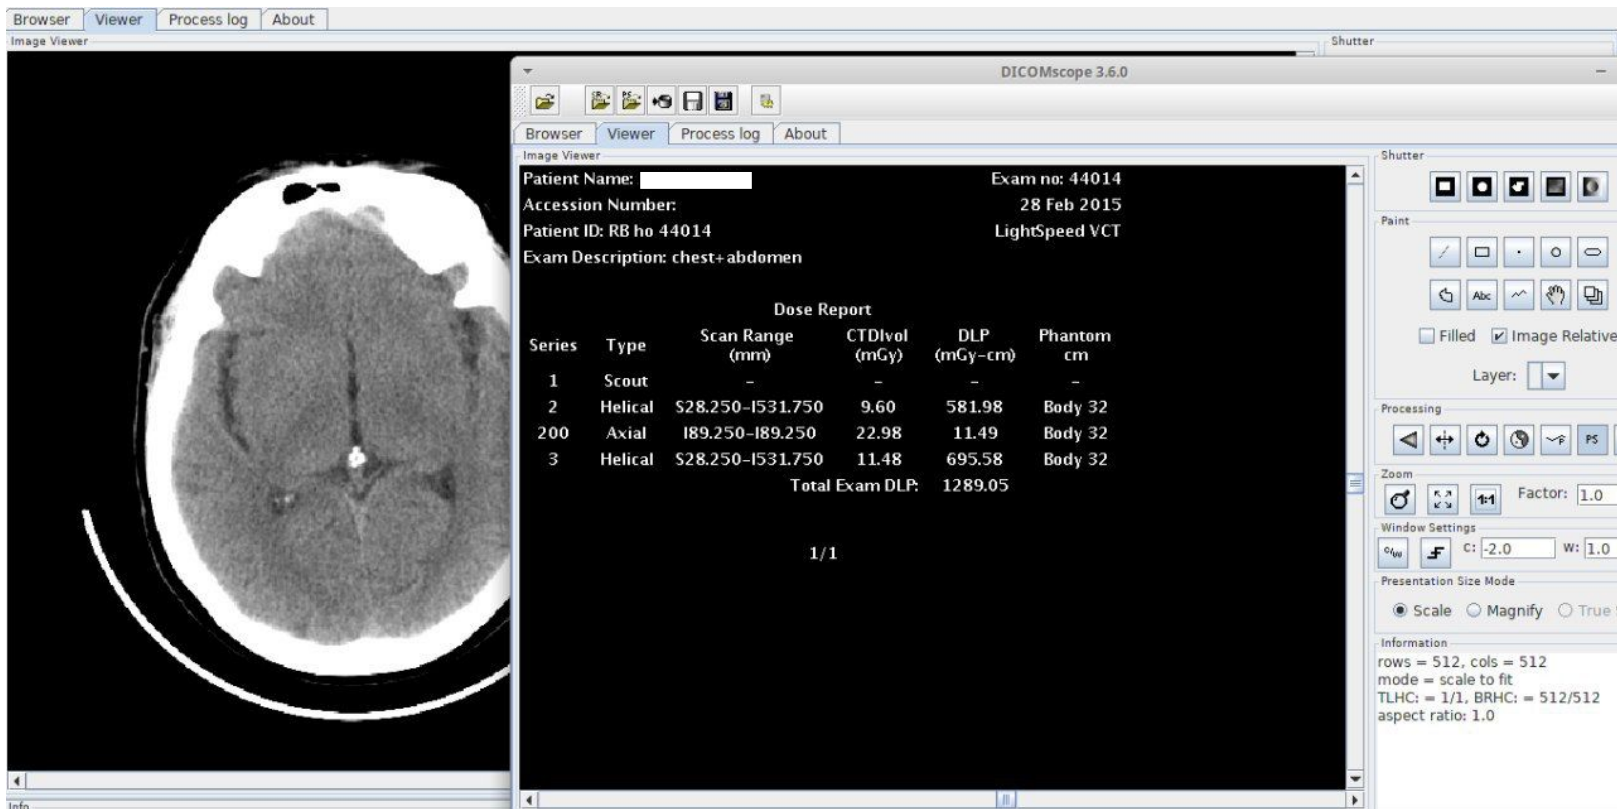
How to retrieve patient docs from PACS?

We can create a dataset that contains any unique patient data (for example, patient name, patient id and so on), and after that, we can get all the related results with C-GET request.

Dataset in DICOM format can be created with `dcmodify` from DCMTK:

```
dcmodify --create-file -i "(0010,0010)=PATIENT_NAME" query_file.dcm
```

How to retrieve patient docs from PACS?



The screenshot displays the DICOMScope 3.6.0 software interface. On the left, a CT scan image of a head is visible. The central panel shows patient information and a dose report table. The right panel contains various toolbars for image manipulation and processing.

Patient Information:

- Patient Name: [REDACTED]
- Exam no: 44014
- Accession Number: 28 Feb 2015
- Patient ID: RB ho 44014
- LightSpeed VCT
- Exam Description: chest+abdomen

Dose Report Table:


Series	Type	Scan Range (mm)	CTDIvol (mGy)	DLP (mGy-cm)	Phantom cm
1	Scout	-	-	-	-
2	Helical	528.250-1531.750	9.60	581.98	Body 32
200	Axial	189.250-189.250	22.98	11.49	Body 32
3	Helical	528.250-1531.750	11.48	695.58	Body 32
Total Exam DLP:				1289.05	

1/1


Right Panel Tools:

- Shutter: [Icons for window management]
- Paint: [Icons for drawing and annotation]
- Processing: [Icons for image processing]
- Zoom: [Icons for zooming, Factor: 1.0]
- Window Settings: [Icons for window settings, C: -2.0, W: 1.0]
- Presentation Size Mode: Scale Magnify True S
- Information: rows = 512, cols = 512, mode = scale to fit, TLHC: = 1/1, BRHC: = 512/512, aspect ratio: 1.0

How to retrieve patient docs from PACS?



Infrastructure Services



Kubeflow

[Documentation](#) / [About](#) / [Kubeflow](#)

Kubeflow

Quickly get running with your ML Workflow

The Kubeflow project is dedicated to making deployments of machine learning (ML) workflows on Kubernetes simple, portable and scalable. Our goal is not to recreate other services, but to provide a straightforward way to deploy best-of-breed open-source systems for ML to diverse infrastructures. Anywhere you are running Kubernetes, you should be able to run Kubeflow.

Kubeflow

Notebook Servers

default
✓ istio-system
kube-public
kube-system
kubeflow


+ NEW SERVER

Status	Name	Created	Image	CPU	Memory	Volumes	Actions
--------	------	---------	-------	-----	--------	---------	---------

More than 300 hosts can be found

Kubeflow


New Notebook Server

 **Name**

Specify the name of the Notebook Server and the Namespace it will belong to.

[Notebook Server's Name](#)
secureserver


[Namespace](#)
istio-system

 **Image**

A starter Jupyter Docker Image with a baseline deployment and typical ML packages.

Standard Custom

[Custom Image](#)
ubuntu:latest

 **CPU**

Specify the total amount of CPU reserved by your Notebook Server. For CPU-intensive workloads, you can choose more than 1 CPU (e.g. 1.5).

[CPU](#)
2.0

Kubeflow

Memory

Specify the total amount of RAM reserved by your Notebook Server (e.g. 2.0Gi).

Memory

4.0Gi

Workspace Volume

Configure the Volume to be mounted as your personal Workspace.


For example, to create an empty Workspace: `New notebook-workspace, 10, /home/jovyan, ReadWriteOnce`







Type	Name	Size (Gi)	Mount Path	Access Mode
New 	secureserver	100	/home/jovyan	ReadWriteOnce 

Data Volumes

Configure the Volumes to be mounted as your Datasets.

For example, to create an empty Data Volume: `New, data-volume-1, 5, /home/jovyan/data-volume-1, ReadWriteOnce`



Type	Name	Size (Gi)	Mount Path	Access Mode	
Existing 	root-value	10	/	ReadWriteOn 	
Existing 	home-value	10	/home/	ReadWriteOn 	

Kubeflow

jupyter Untitled Last Checkpoint: 10/06/2019 (autosaved)



File Edit View Insert Cell Kernel Widgets Help

Trusted | Python 3



```
In [8]: from __future__ import print_function

import tensorflow as tf
from tensorflow import keras

# Helper libraries
import numpy as np
import os
import subprocess
import argparse

# Reduce spam logs from s3 client
os.environ['TF_CPP_MIN_LOG_LEVEL']='3'

def preprocessing():
    fashion_mnist = keras.datasets.fashion_mnist
    (train_images, train_labels), (test_images, test_labels) = fashion_mnist.load_data()

    # scale the values to 0.0 to 1.0
    train_images = train_images / 255.0
    test_images = test_images / 255.0

    # reshape for feeding into the model
    train_images = train_images.reshape(train_images.shape[0], 28, 28, 1)
    test_images = test_images.reshape(test_images.shape[0], 28, 28, 1)
```



Kubeflow

jupyter Untitled Last Checkpoint: 10/06/2019 (autosaved)



jupyter

```
$ ls
bin boot dev etc home lib lib64 media mnt opt proc root run sbin srv sys tf tmp usr var
$ uname -a
Linux myjupyter-0 4.14.146-119.123.amzn2.x86_64 #1 SMP Mon Sep 23 16:58:43 UTC 2019 x86_64 x86_64 x86_64 GNU/Linux
$ lsb_release -a
No LSB modules are available.
Distributor ID: Ubuntu
Description:    Ubuntu 18.04.2 LTS
Release:        18.04
Codename:       bionic
$
```

```
(train_images, train_labels), (test_images, test_labels) = fashion_mnist.load_data()
```

```
# scale the values to 0.0 to 1.0
```

```
train_images = train_images / 255.0
```

```
test_images = test_images / 255.0
```

```
# reshape for feeding into the model
```

```
train_images = train_images.reshape(train_images.shape[0], 28, 28, 1)
```

```
test_images = test_images.reshape(test_images.shape[0], 28, 28, 1)
```



Thank you for attention!

Any questions?



 @dnkolegov

 @manmoleculo

 github.com/sdnewhop/grinder

 sdnewhop.github.io/AISec/