

# OSC - UAV Use Case Guideline

☰ Tags	技術 概念
№ ID	WD-72
🕒 Last edited time	@February 6, 2024 8:51 PM

## Introduction

We have developed a new use case for OSC, successfully generating a new model for UAV path prediction by incorporating UAV data and modifying the pipeline.

### Introduction

#### Step 1. Add UAV data in InfluxDB

##### 1-1 Create a new bucket

##### 1-2 Copy and revise recipe file `RECIPE_EXAMPLE/example_recipe_latest_stable.yaml`

##### 1-3 Accessing applications in the cluster using port forwarding to send data

#### Step 2. Create a pipeline

##### 2-1 create a new pipeline

#### Step 3. Creating training job

##### 3-1 Creating an new training job

##### 3-2 Data extraction

#### Step 4. Test predictions using model deployed on Kserve

##### 4-1 Deploy trained UAV prediction model on Kserve

##### 4-2 (problem) Failed calling webhook "inferenceservice.kserve-webhook-server.default"

##### 4-3 Test predictions using model deployed on Kserve

## Step 1. Add UAV data in InfluxDB

### ▼ 1-1 Create a new bucket

- find influxdb pod

```
kubectl get pods -A
```



```
I have no name!@my-release-influxdb-5b77fc46b4-5f6f7:/# influx bucket list --org primary --token VJpoNpqeVnjzvhpPm8jZ
ID                               Name           Retention      Shard group duration  Organization ID  Schema Type
843ee9b113635d9d                UAVData        infinite       168h0m0s              103894585d415659  implicit
4d219163d016dbfb                UEData         infinite       168h0m0s              103894585d415659  implicit
32cb15b323ef57cf                _monitoring    168h0m0s      24h0m0s               103894585d415659  implicit
873e1b5d0ea6c982                _tasks         72h0m0s       24h0m0s               103894585d415659  implicit
7f4bff75d6adf05d                primary        infinite       168h0m0s              103894585d415659  implicit
```

## ▼ 1-2 Copy and revise recipe

file `RECIPE_EXAMPLE/example_recipe_latest_stable.yaml`

- Newfile\_name : `example_recipe_latest_stable_copy.yaml`
- change bucket of datalake

```
traininghost:
  ip_address: 192.168.190.140

datalake:
  influxdb:
    host: 192.168.190.140
    port: 8086
    orgname: primary
    bucket: UAVData
    token: VJpoNpqeVnjzvhpPm8jZ
```

- Once updated, follow the below steps for reinstall of some components

```
bin/uninstall.sh
bin/install.sh -f RECIPE_EXAMPLE/example_recipe_latest_stable.yaml
```

## ▼ 1-3 Accessing applications in the cluster using port forwarding to send data

- revise `insert.py`
- File\_path : `aimlfw/qp/qp/insert.py`

- Newfile\_file : Mitlab\_insert.py

```
import pandas as pd
from influxdb_client import InfluxDBClient
from influxdb_client.client.write_api import SYNCHRONOUS
import datetime

class INSERTDATA:
    def __init__(self):
        self.client = InfluxDBClient(url = "http://192.168

def explode(df):
    for col in df.columns:
        if isinstance(df.iloc[0][col], list):
            df = df.explode(col)
        d = df[col].apply(pd.Series)
        df[d.columns] = d
        df = df.drop(col, axis=1)
    return df

def jsonToTable(df):
    df.index = range(len(df))
    cols = [col for col in df.columns if isinstance(df.il
    if len(cols) == 0:
        return df
    for col in cols:
        d = explode(pd.DataFrame(df[col], columns=[co
        d = d.dropna(axis=1, how='all')
        df = pd.concat([df, d], axis=1)
        df = df.drop(col, axis=1).dropna()
    return jsonToTable(df)

def data_normalization(df):
    import numpy as np
```

```

df.drop('time', axis=1, inplace=True)
# Normalization
std = []
mean = []

# Iterate over columns
for col in df.columns:
    std_val = df[col].std()
    mean_val = df[col].mean()
    std.append(std_val)
    mean.append(mean_val)
    df[col] = (df[col] - mean_val) / std_val
return df

def time(df):
    df.index = pd.date_range(start=datetime.datetime.now(
df['time'] = df['time'].apply(lambda x: str(x))
    return df

def populatedb():
    df = pd.read_csv('/home/mitlab/osc/aimlfw-dep/qp/qp/o
df = jsonToTable(df)
df = time(df)
df = data_normalization(df)
print(df)
db = INSERTDATA()
write_api = db.client.write_api(write_options=SYNCHRO
write_api.write(bucket="UAVData",record=df, data_fram

populatedb()

```

```

from influxdb_client import InfluxDBClient
from influxdb_client.client.write_api import SYNCHRONOUS
import datetime

# num=0

class INSERTDATA:
    def __init__(self):
        self.client = InfluxDBClient(url = "http://192.168.190.140:8086", token="VJpoNpqeVnjzvhpPm8jZ")

def explode(df):
    for col in df.columns:
        if isinstance(df.iloc[0][col], list):
            df = df.explode(col)
            d = df[col].apply(pd.Series)
            df[d.columns] = d
            df = df.drop(col, axis=1)
    return df

def jsonToTable(df):
    df.index = range(len(df))
    cols = [col for col in df.columns if isinstance(df.iloc[0][col], dict) or isinstance(df.iloc[0][col], list)]
    if len(cols) == 0:
        return df
    for col in cols:
        d = explode(pd.DataFrame(df[col], columns=[col]))
        d = d.dropna(axis=1, how='all')
        df = pd.concat([df, d], axis=1)
        df = df.drop(col, axis=1).dropna()
    return jsonToTable(df)

def data_normalization(df):
    import numpy as np
    df.drop('time', axis=1, inplace=True)
    # Normalization
    std = []
    mean = []

    # Iterate over columns
    for col in df.columns:
        std_val = df[col].std()
        mean_val = df[col].mean()
        std.append(std_val)
        mean.append(mean_val)
        df[col] = (df[col] - mean_val) / std_val
    return df

def time(df):
    df.index = pd.date_range(start=datetime.datetime.now(), freq='10ms', periods=len(df))
    df['time'] = df['time'].apply(lambda x: str(x))
    return df

def populatedb():
    df = pd.read_csv('/home/mitlab/osc/aimlfw-dep/qp/qp/original_dataset.csv')
    df = jsonToTable(df)
    df = time(df)
    df = data_normalization(df)
    print(df)
    db = INSERTDATA()
    write_api = db.client.write_api(write_options=SYNCHRONOUS)
    write_api.write(bucket="UAVData", record=df, data_frame_measurement_name="liveCell", org="primary")

populatedb()

```

- Follow below command to port forward to access Influx DB

- **Step 1 : Check influx service name and port**

```
kubectl get service -A
```

NAMESPACE	AGE	NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)
default	97m	kubernetes	ClusterIP	10.96.0.1	<none>	443/TCP
default	53m	my-release-influxdb	ClusterIP	10.104.251.208	<none>	8086/TCP,8088/TCP
kube-system	97m	kube-dns	ClusterIP	10.96.0.10	<none>	53/UDP,53/TCP,9153/TCP
kubeflow	87m	cache-server	ClusterIP	10.104.106.133	<none>	443/TCP

- **Step 2 : Open new terminal and follow below command to port forward to Influx DB**

```
kubectl port-forward svc/<Your influxDB service name> 8
```

For example :

```
kubectl port-forward svc/my-release-influxdb 8086:8086
```

- **Step 3 : Back to the terminal and run this command to insert data**

```
python3 Mitlab_insert.py
```

- To check inserted data in Influx DB, **execute below command inside the Influx DB container:**

- **Step 1. Get into influxdb pod**

```
kubectl exec -it my-release-influxdb-5b77fc46b4-5f6f7 -
```

- **Step 2. Check the data in the container**

```
influx query 'from(bucket: "UAVData") |> range(start:
```

For example:

```
influx query 'from(bucket: "UAVData") |> range(start:
```

and you will see the information like this figure.

2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.63366000Z	-0.2320283645053262
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.64366000Z	-0.25532153882945974
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.65366000Z	-0.25764142343422414
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.66366000Z	-0.054079252752072635
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.67366000Z	-0.25605171347674366
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.68366000Z	-0.14234969094599112
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.69366000Z	-0.2301903917534637
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.70366000Z	-0.11034514759404551
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.71366000Z	-0.23063466925392485
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.72366000Z	-0.12823137618472721
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.73366000Z	-0.16264210613445496
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.74366000Z	-0.1663222910353533
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.75366000Z	-0.161013695811879
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.76366000Z	-0.16814310639508073
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.77366000Z	-0.16775908622490958
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.78366000Z	-0.16367271291196317
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.79366000Z	-0.15166757775498024
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.80366000Z	-0.152450088358923
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.81366000Z	-0.1514870783374395
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.82366000Z	-0.15812570611815088
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.83366000Z	-0.1507198819204624
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.84366000Z	-0.13701048521460363
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.85366000Z	-0.16488086617798733
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.86366000Z	-0.1703062925518383
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.87366000Z	-0.13073826229186844
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.88366000Z	-0.13085231317784763
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.89366000Z	-0.129645380655014
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.90366000Z	-0.13520196208090348
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.91366000Z	-0.24196579954471686
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.92366000Z	-0.261745630224182
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.93366000Z	-0.1703419495760048
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.94366000Z	-0.1793182525770819
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.95366000Z	-0.17734446594781345
2021-03-10T10:49:07.232011624Z	2023-12-05T10:49:07.232011624Z	z	liveCell	2023-12-04T18:59:18.96366000Z	-0.1893860867563316

## Step 2. Create a pipeline

### ▼ 2-1 create a new pipeline

- Revise OSC qoe.iphy
- File\_path : connect 32088 port
- Newfile\_path : Mitlab\_pipeline.iphy

```
import kfp
import kfp.components as components
import kfp.dsl as dsl
from kfp.components import InputPath, OutputPath

def train_export_model(trainingjobName: str, epochs: str,

import tensorflow as tf
```



```

from numpy import array
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Dense, Flatten, Dr
import numpy as np
print("numpy version")
print(np.__version__)
import pandas as pd
import os
from featurestoresdk.feature_store_sdk import FeaturesS
from modelmetricsdk.model_metrics_sdk import ModelMetr

import requests
from tensorflow import keras
from sklearn.model_selection import train_test_split

fs_sdk = FeatureStoreSdk()
mm_sdk = ModelMetricsSdk()

features = fs_sdk.get_features(trainingjobName, ['x',
features = features.astype('float64')
print("Dataframe:")
print(features)
print(features.dtypes)

features = features.to_numpy()
print("features.dtype:")
print(features.dtype)

one_set_data_count = 2500
set_count = 1
input_data_series_count = 100 # hyper parameter
separate_data = []
series = []
label = []
for i in range(set_count):
    separate_data.append(features[i*one_set_data_count

```

```

        for j in range(one_set_data_count-input_data_series.append(separate_data[i][j:j+input_data_label.append(separate_data[i][j+input_data_ser
separate_data = np.array(separate_data)
series = np.array(series)
label = np.array(label)

# make train and test
x_train, x_test, y_train, y_test = train_test_split(se
x_test, x_validate, y_test, y_validate = train_test_sp
print("x_train:", x_train.dtype)
print("y_train:", y_train.dtype)
print("x_validate:", x_validate.dtype)
print("y_validate:", y_validate.dtype)

# Define the model architecture
model = Sequential()
model.add(LSTM(units=64, input_shape=(input_data_serie
model.add(Dropout(0.3))
model.add(Dense(32, activation='relu'))
model.add(Dense(3, activation='linear'))

# Compile the model
model.compile(loss='mean_squared_error', optimizer='ad

# Train the model on your data
batch_size = 4
model.fit(x_train, y_train, batch_size=batch_size, epo

model.save("./")
mm_sdk.upload_model("./", trainingjobName, version)

BASE_IMAGE = "traininghost/pipelineimage:latest"

def train_and_export(trainingjobName: str, epochs: str, ve

```

```

trainOp = components.func_to_container_op(train_export
# Below line to disable caching of pipeline step
trainOp.execution_options.caching_strategy.max_cache_s
trainOp.container.set_image_pull_policy("IfNotPresent"

@dsl.pipeline(
    name='MITLAB_training_pipeline',
    description='UAVData_pipeline'
)
def super_model_pipeline(
    trainingjob_name: str, epochs: str, version: str):

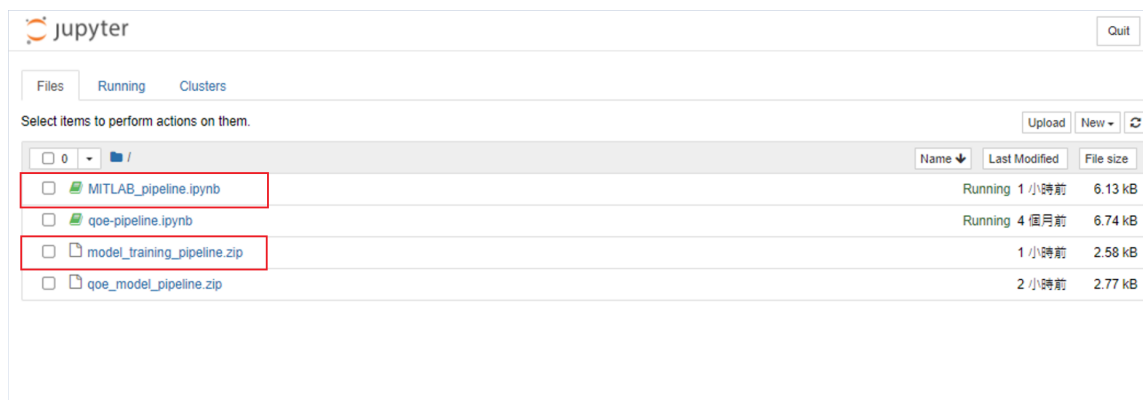
    train_and_export(trainingjob_name, epochs, version)

pipeline_func = super_model_pipeline
file_name = "model_training_pipeline"
kfp.compiler.Compiler().compile(pipeline_func, '{}.zip'.fo

import requests
pipeline_name="MITLAB_training_pipeline"
pipeline_file = file_name+'.zip'
requests.post("http://tm.traininghost:32002/pipelines/{}"/u

```

- Produce a new pipeline



## Step 3. Creating training job

### ▼ 3-1 Creating an new training job

- Create an new training job on aiml-dashboard

The screenshot shows the 'AI/ML Management Dashboard' with a 'Training Jobs' dropdown menu open. The form contains the following fields:

- Training Job Name\***: Text input field.
- Training Function\***: Dropdown menu with the text '--- Select Training Function ---'.
- Experiment Name\***: Dropdown menu with the text '--- Select Experiment ---'.
- Datalake Source\***: Dropdown menu with the text '--- Select Datalake Source ---'.
- Feature Name\***: Text input field.
- Feature Filter**: Text input field.
- Hyper Parameters**: Text input field.
- Enable versioning**
- Description**: Text input field.

A blue button labeled 'Create Training Job' is located at the bottom of the form.

- Use the default parameter

Parameter	Value
Training Job Name	test1208
Training Function	model_training_pipeline
Training Function Version Name	1
Experiment Name	Default
Datalake Source	Influx DB
_measurement	liveCell
bucket	UAVData
Feature Name	*

Feature Filter	
Hyper Parameters	epochs:1
Description	test

- training job info

### Training Job Info ×

---

Training Job Name  
test1208

Version  
1

Description  
test

Feature Names  
\*

Training Function Name  
MITLAB\_training\_pipeline

Experiment Name  
Default

Feature Filter

Hyper Parameters  
epochs:1,trainingjob\_name:test1208

### Metrics

No data available

Enable versioning

Training Function Version  
1

Datalake Source  
Influx DB

\_measurement  
liveCell

bucket  
UAVData

Model URL  
http://192.168.190.140:32002/model/test1208/1/Model.zip

## ▼ 3-2 Data extraction

- **(Problem) The module cannot successfully download in the data exaction pod.**
  - Data extraction pod error message (**CoreDNS Problem**)

```

:: resolving dependencies :: org.apache.spark#spark-submit-parent-eabfd46-6597-4711-986e-646d1ccda88;1.0
  confs: [default]
You probably access the destination server through a proxy server that is not well configured.
You probably access the destination server through a proxy server that is not well configured.
You probably access the destination server through a proxy server that is not well configured.
You probably access the destination server through a proxy server that is not well configured.
:: resolution report :: resolve 79ms :: artifacts dl 0ms
  :: modules in use:
  |-----|-----|
  |   conf   |   modules   |   artifacts   |
  | number| search|download|evicted| number|download| |
|---|---|---|---|---|---|---|
  | default | 1 | 0 | 0 | 0 | 0 | 0 |
  |-----|-----|

:: problems summary ::
::: WARNINGS
Host repo1.maven.org not found. url=https://repo1.maven.org/maven2/com/datastax/spark/spark-cassandra-connector_2.12/3.0.1/spark-cassandra-connector_2.12-3.0.1.pom
Host repo1.maven.org not found. url=https://repo1.maven.org/maven2/com/datastax/spark/spark-cassandra-connector_2.12/3.0.1/spark-cassandra-connector_2.12-3.0.1.jar
Host repos.spark-packages.org not found. url=https://repos.spark-packages.org/com/datastax/spark/spark-cassandra-connector_2.12/3.0.1/spark-cassandra-connector_2.12-3.0.1.pom

```

- o To resolve **CoreDNS Problem** in kubernetes:
  - Step 1. Enter the data extraction pod and **add nameserver 8.8.8.8**(Google's DNS server) to `/etc/resolv.conf` in the pod ,restart the data extraction pod and restart the training job again to download the essential module.

```
kubectl exec -it --namespace=traininghost data-extra
```

```

cat << EOF > /etc/resolv.conf
nameserver 8.8.8.8
nameserver 10.96.0.10
search traininghost.svc.cluster.local svc.cluster.local
options ndots:5
EOF

```

check the modules already downloaded or not.

```

:: resolution report :: resolve 29375ms :: artifacts dl 4210ms
:: modules in use:
com.datastax.oss#java-driver-core-shaded;4.10.0 from central in [default]
com.datastax.oss#java-driver-mapper-runtime;4.10.0 from central in [default]
com.datastax.oss#java-driver-query-builder;4.10.0 from central in [default]
com.datastax.oss#java-driver-shaded-guava;25.1-jre-graal-sub-1 from central in [default]
com.datastax.oss#native-protocol;1.4.12 from central in [default]
com.datastax.spark#spark-cassandra-connector-driver-2.12;3.0.1 from central in [default]
com.datastax.spark#spark-cassandra-connector-2.12;3.0.1 from central in [default]
com.github.spotbugs#spotbugs-annotations;3.1.12 from central in [default]
com.github.stephenc.jcip#jcip-annotations;1.0-1 from central in [default]
com.google.code.findbugs#jsr305;3.0.2 from central in [default]
com.thoughtworks.paranamer#paranamer;2.8 from central in [default]
com.typesafe#config;1.4.1 from central in [default]
io.dropwizard.metrics#metrics-core;4.1.16 from central in [default]
org.apache.commons#commons-lang3;3.9 from central in [default]
org.hdrhistogram#HdrHistogram;2.1.12 from central in [default]
org.reactivestreams#reactive-streams;1.0.3 from central in [default]
org.scala-lang#scala-reflect;2.12.11 from central in [default]
org.slf4j#slf4j-api;1.7.26 from central in [default]
-----
|           | modules | artifacts |
|   conf   | number | search | dl | dl | evicted | | number | dl | dl |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| default  | 18     | 18     | 18 | 0 | 0       | | 18     | 18 | 0 |
-----
:: retrieving :: org.apache.spark#spark-submit-parent-af6d378f-b9c5-4b2a-90cc-0e5f12c5162b
confs: [default]
18 artifacts copied, 0 already retrieved (17880KB/28ms)
23/12/05 11:26:46 WARN NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes where applicable
Using Spark's default log4j profile: org/apache/spark/log4j-defaults.properties
Setting default log level to "WARN".
To adjust logging level use sc.setLogLevel(newLevel). For SparkR, use setLogLevel(newLevel).
2023-12-05 11:26:51,522 | _internal.py 225 _log() | INFO | 172.16.189.238 - - [05/Dec/2023 11:26:51] "GET /task-status/testdate1205 HTTP/1.1" 200 -
2023-12-05 11:26:54,766 | Pipeline.py 43 load_data() | INFO | Source:<source.InfluxSource.InfluxSource object at 0x7f93e80ae828>
2023-12-05 11:26:54,767 | InfluxSource.py 78 load() | DEBUG | Started Data Extraction for Influx Source from(bucket:"UAVData") |> range(start: 0, stop
ey:["time"], columnkey: ["_field"], valueColumn: "_value")
2023-12-05 11:27:01,566 | _internal.py 225 _log() | INFO | 172.16.189.238 - - [05/Dec/2023 11:27:01] "GET /task-status/testdate1205 HTTP/1.1" 200 -
2023-12-05 11:27:11,623 | _internal.py 225 _log() | INFO | 172.16.189.238 - - [05/Dec/2023 11:27:11] "GET /task-status/testdate1205 HTTP/1.1" 200 -
2023-12-05 11:27:21,652 | _internal.py 225 _log() | INFO | 172.16.189.238 - - [05/Dec/2023 11:27:21] "GET /task-status/testdate1205 HTTP/1.1" 200 -
2023-12-05 11:27:31,709 | _internal.py 225 _log() | INFO | 172.16.189.238 - - [05/Dec/2023 11:27:31] "GET /task-status/testdate1205 HTTP/1.1" 200 -
2023-12-05 11:27:41,823 | _internal.py 225 _log() | INFO | 172.16.189.238 - - [05/Dec/2023 11:27:41] "GET /task-status/testdate1205 HTTP/1.1" 200 -

```

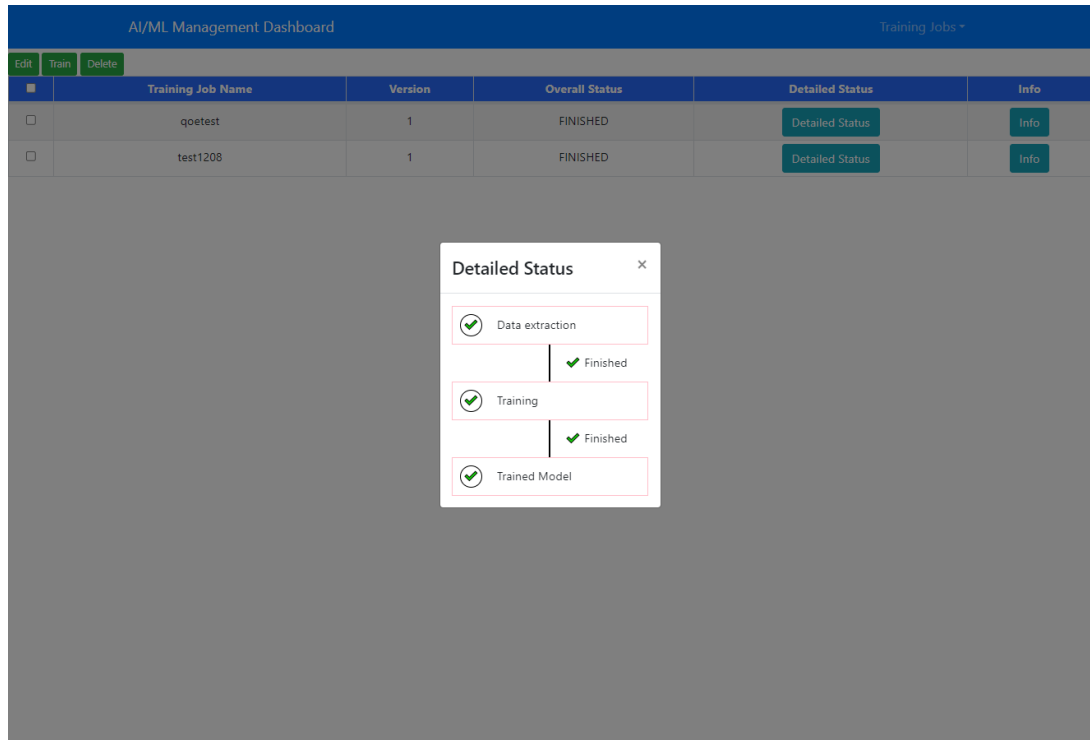
- Step 2. After the pod successfully downloads the module, enter the data extraction pod and **restore /etc/resolv.conf**.

```

cat << EOF > /etc/resolv.conf
nameserver 10.96.0.10
search traininghost.svc.cluster.local svc.cluster.lo
options ndots:5
EOF

```

- Re-execute the training job, wait for minutes then the model is complete.



## Step 4. Test predictions using model deployed on Kserve

### ▼ 4-1 Deploy trained UAV prediction model on Kserve

- Create `uav.yaml` file with below contents

```

apiVersion: "serving.kserve.io/v1beta1"
kind: "InferenceService"
metadata:
  name: uav-model
spec:
  predictor:
    tensorflow:
      storageUri: "http://192.168.190.140:32002/model/test"
      runtimeVersion: "2.5.1"
    resources:
      requests:
        cpu: 1
        memory: 0.5Gi

```



```
limits:
  cpu: 2
  memory: 0.5Gi
```

- deploy uav-model

```
kubectl apply -f uav.yaml -n kserve-test
```

- Check running state of pod

```
kubectl get pods -n kserve-test
```

```
root@mitlab-virtual-machine:/home/mitlab/osc/aimlfw-dep# kubectl apply -f uav.yaml -n kserve-test
inferenceservice.serving.kserve.io/uav-model configured
root@mitlab-virtual-machine:/home/mitlab/osc/aimlfw-dep# kubectl get pods -n kserve-test
```

NAME	READY	STATUS	RESTARTS	AGE
qoe-model-predictor-default-00001-deployment-68d85bf59b-45j4g	2/2	Running	0	124d
uav-model-predictor-default-00001-deployment-69d4f54ddf-18nhd	1/2	Running	0	2s

#### ▼ 4-2 (problem) Failed calling webhook "inferenceservice.kserve-webhook-server.defaulter"

- deleting 2 webhook config

```
kubectl delete mutatingwebhookconfigurations.admissionregi
kubectl delete validatingwebhookconfigurations.admissionre
```

- inferenceservice.serving.kserve.io/uav-model configured

```
kubectl apply -f uav.yaml -n kserve-test
```

#### ▼ 4-3 Test predictions using model deployed on Kserve

- Use below command to obtain Ingress port for Kserve.

```
kubectl get svc istio-ingressgateway -n istio-system
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
istio-ingressgateway	LoadBalancer	10.101.170.189	<pending>	15021:32140/TCP,80:32576 TCP,443:32435/TCP,15012:32114/TCP,15443:31866/TCP	33m

- Create `uav.sh` file with following contents

```
nano uav.sh
```

- Copy the below content and update the “**IP of host**” where Kserve is deployed and ingress “**port**” of Kserve obtained using above method.

```
model_name=uav-model  
curl -v -H "Host: $model_name.kserve-test.example.com" htt
```

For example:

```
model_name=uav-model  
curl -v -H "Host: $model_name.kserve-test.example.com" htt
```

- After complete update, create sample data for predictions in file **input\_uav.json**.

```
nano input_uav.json
```

Add the following content in `input_uav.json` file.

```
{"signature_name": "serving_default", "instances": [[[0.74  
[0.9348491377517286, 0.5797269721343681, 0.15480809  
[0.32685562293198867, 0.40325520693639316, 0.379231
```

- Use command below to trigger predictions.

```
source uav.sh
```