

*“Houston,
we have a model...”*

Introduction to MLOps

Rui Quintino

AI & Analytics Strategy Advisor
DevScope

20210219




dev>scope

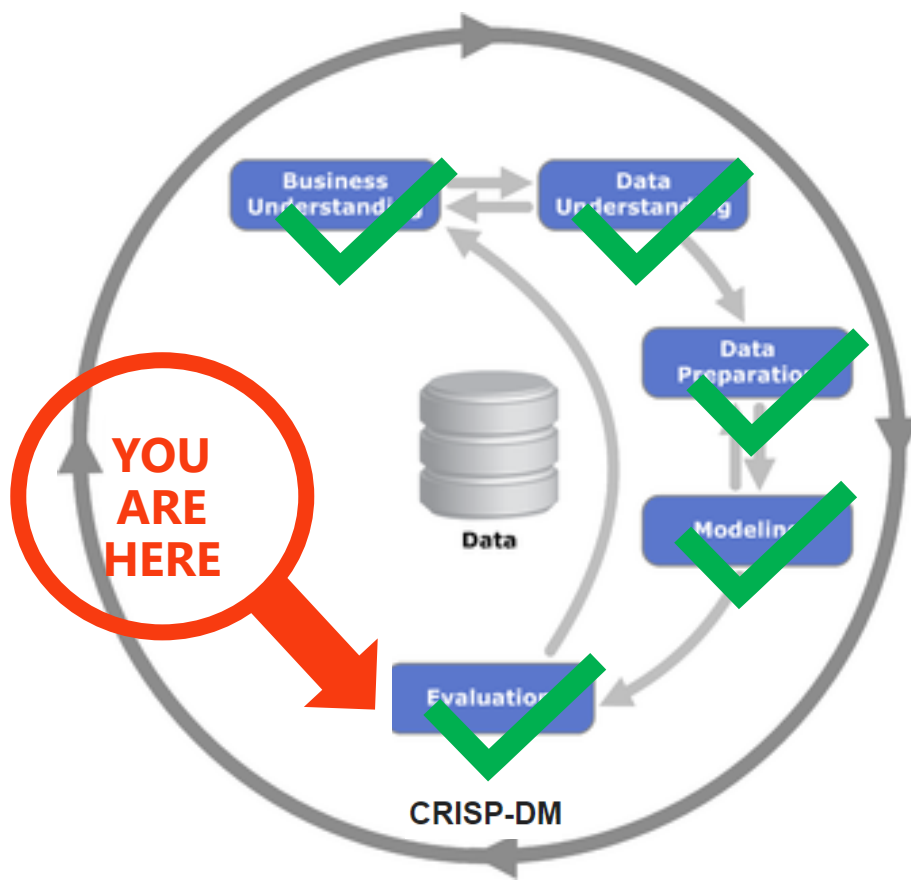
AI & Analytics

Strategy Advisor

 ruí.quintino@devscope.net

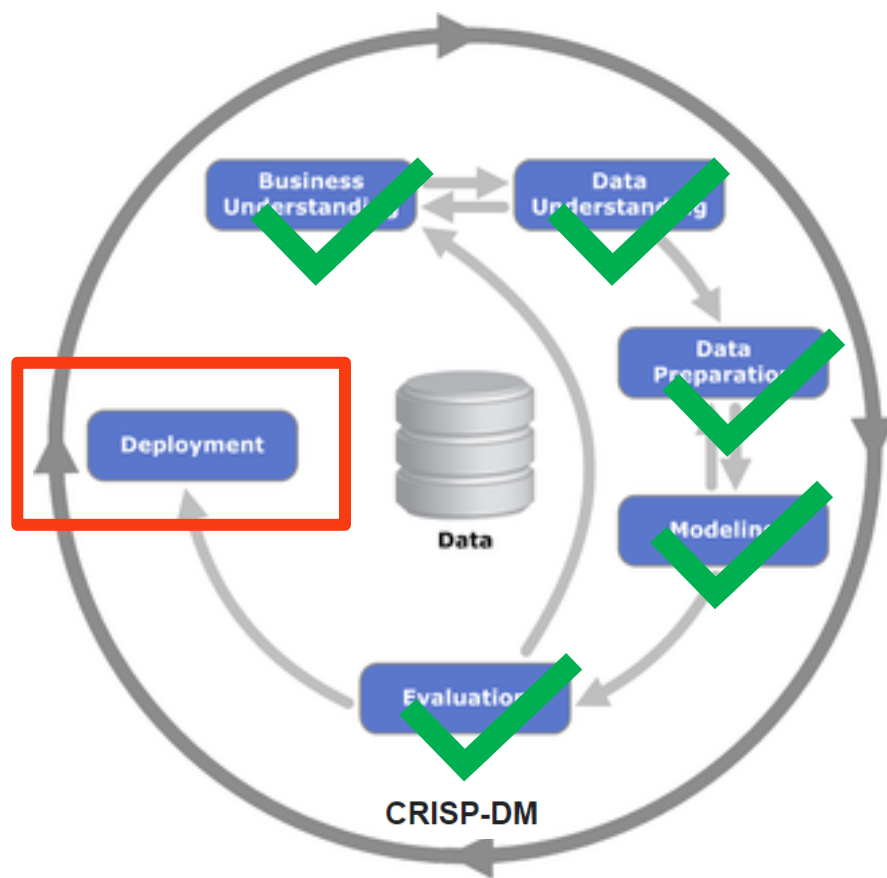
 [Linkedin.com/ in/rquintino](https://www.linkedin.com/in/rquintino)





We have a model... It's great!

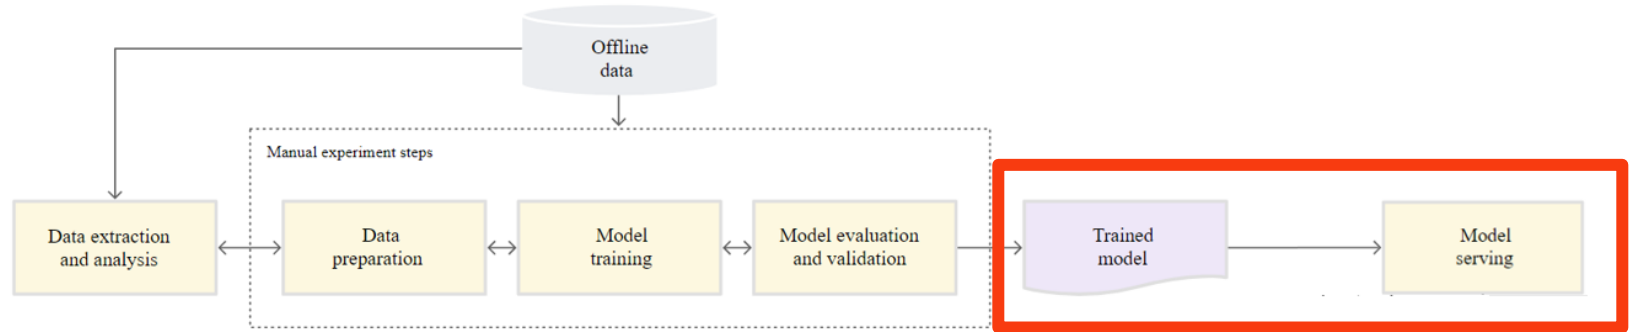
Now what?



We just deploy it!

```
model file name = "model.pkl"
```

```
joblib.dump(svm_model, os.path.join('outputs', model_file_name))
```



Adapted from: <https://cloud.google.com/solutions/machine-learning/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning>

Easy right?

Contents

- 1 Run
- 2 Baseline
- 3 Read data
 - 3.1 Join both so we can calculate sl
 - 4 Corrupted files (Ex: large white zon)
- 5 Predict Next steps? (Forecast)
 - 5.1 Add future/prev labels
- 6 Train/val
 - 6.1 Checks
- 7 Train/Val Splits (KFold)
 - 7.1 Generate & Assign KFolds (Gro
 - 7.2 Check distribution of %samples
 - 7.3 Attribute folds
 - 7.4 Pick Fold & attribute percentage
 - 7.5 Now with a specific perc per stc
- 8 Model pipeline
 - 8.1 Image preproc
 - 8.2 Torch Loader & Transforms
- 9 Model
 - 9.1 Model Config
- 10 Train
- 11 Save Model
- 12 Score
 - 12.1 Load best checkpoint
 - 12.2 Score validation fold
 - 12.3 Score train phase for val fold
 - 12.4 Score Test images (final subm
 - 12.5 Score test storms train images

```
def stack_images(row):
    row_name=fill_nas(row, "file_name", "")
    row_time=fill_nas(row, "diff_time", 0).round(1)

    img1=read_img(row_name["file_name"], row.get("rotate"), row.get("crop"))
    img2=read_img(row_name["file_name_prev2"], row.get("rotate"), row.get("crop"))
    img3=read_img(row_name["file_name_prev4"], row.get("rotate"), row.get("crop"))


    array=(np.moveaxis(np.stack([img1, img2, img3]), 0, -1))

    # Inject two small squares to encode ocean and relative time info
    # Use max time of ~20 days

    max_time=20*3600*24
    time_marker=np.min([255, 150+((row["relative_time"]/max_time)*255).astype("uint8")])
    array[121:140, 121:140, 2]=time_marker
    array[141:160, 141:160, 1]=row["ocean"]*50

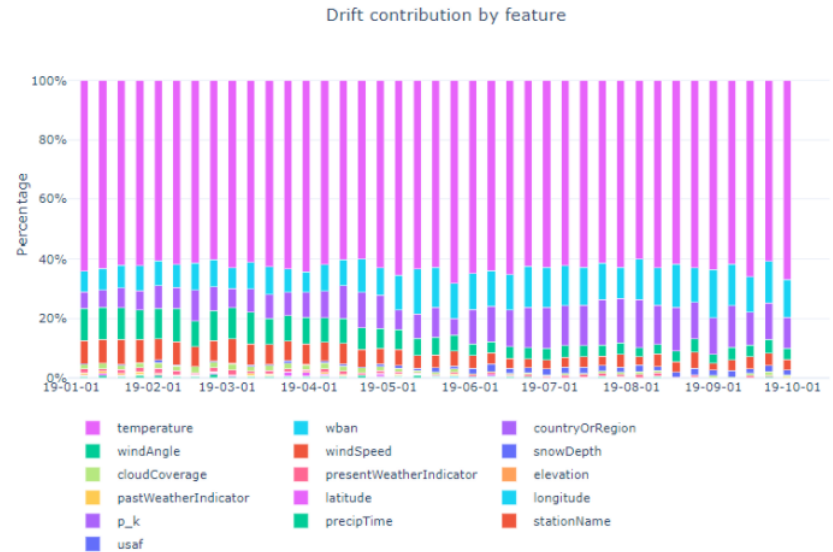
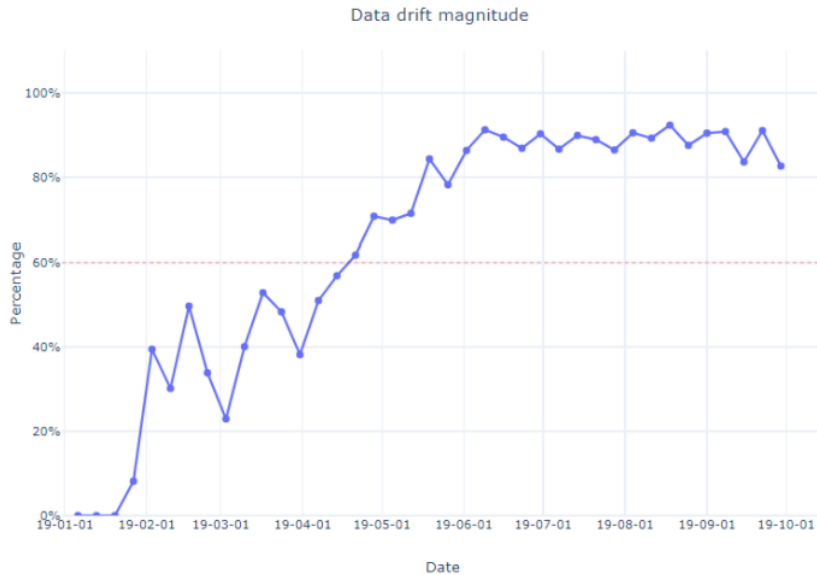
    return pil_image.fromarray(array)
```

executed in 344ms, finished 14:20:49 2021-02-18



What about the data Pipeline/preproc?

Drift overview



<https://github.com/Azure/MachineLearningNotebooks/blob/master/how-to-use-azureml/work-with-data/datadrift-tutorial/datadrift-tutorial.ipynb>

Or Data Drift -> causing Model Performance degradation?

Plus...

- How/who can we retrain if needed (ex: data drift)?
- What are the software requirements?
- How can we later update lib/os versions? (ex: security)
- How can we monitor model performance in production?
- How do we ensure data checks in production?
- How was the model trained? What was the performance?
- How to scale? And Reproduce?
- ...



ML
code

Model & ML code just tip of the iceberg...

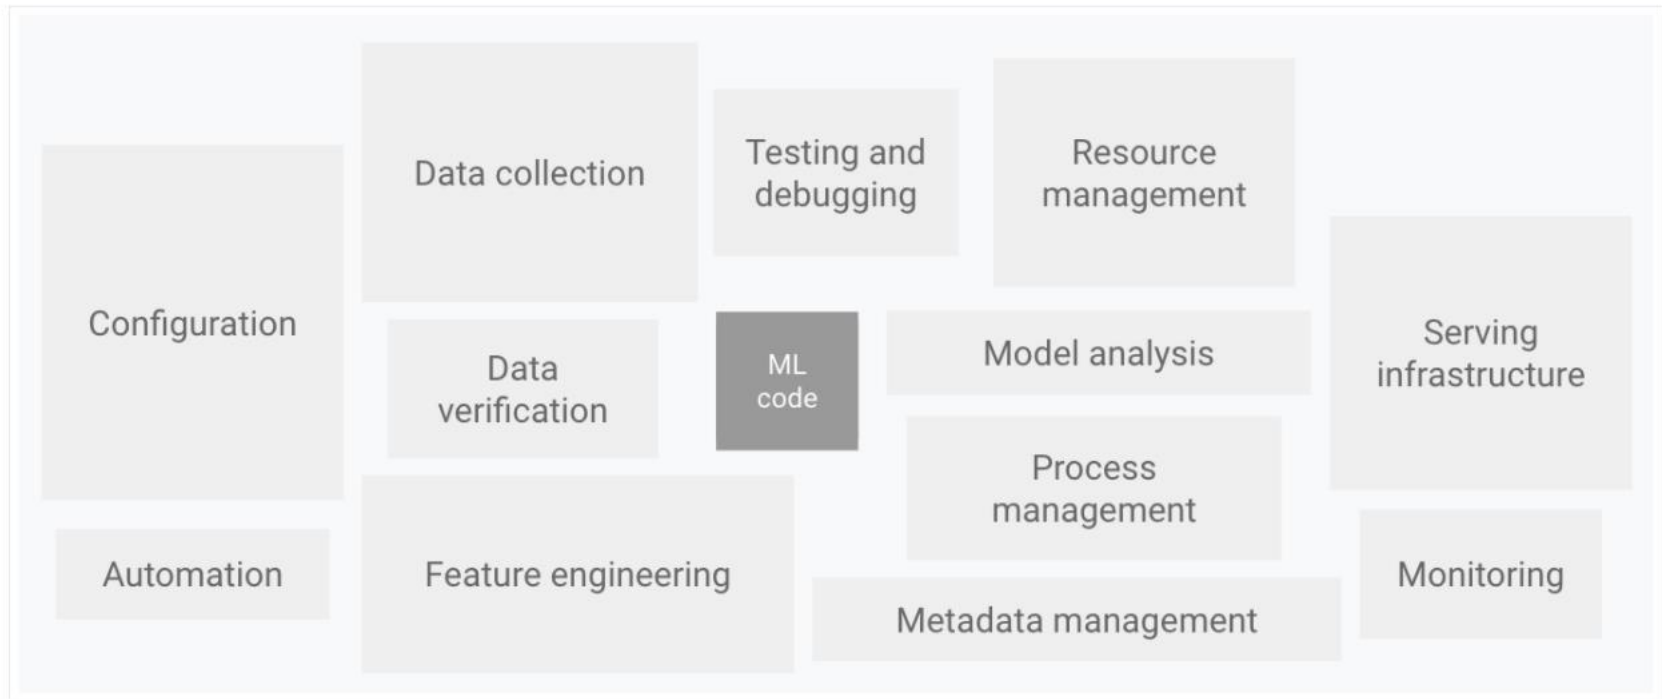


Figure 1. Elements for ML systems. Adapted from [Hidden Technical Debt in Machine Learning Systems](#).

<https://cloud.google.com/solutions/machine-learning/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning>

Model & ML code just tip of the iceberg...

Introducing MLOps – Machine Learning & Operations

DevOps



Code reproducibility



Code testing



App deployment

MLOps



Model reproducibility



Model validation




Model deployment

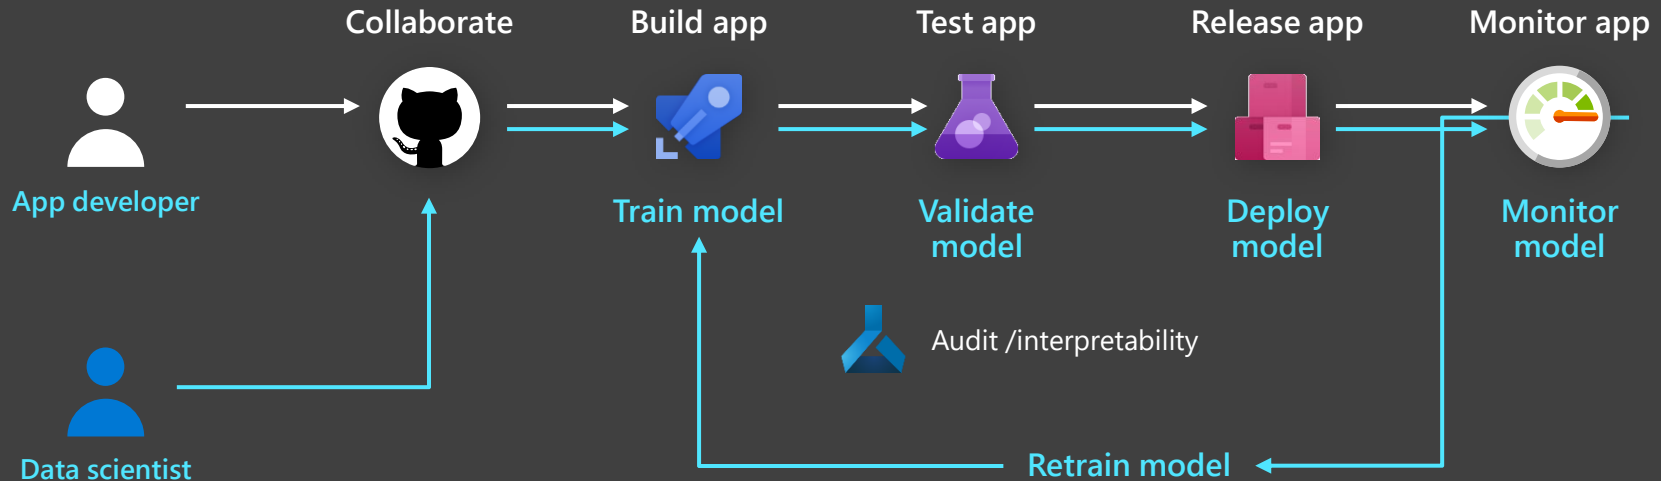


Model retraining

DevOps



MLOps with Azure Machine Learning



✓ Model reproducibility

✓ Model validation

✓ Model deployment

✓ Model retraining



MLOps Demo

GitHub & Azure ML

rquintino / mlops_azureml Private

Unwatch 1 Star 0 Fork 0

generated from machine-learning-apps/ml-template-azure

Code Issues machine-l.../ml-template-azure Template Security Insights Settings

master Python 47 stars 28 forks Updated on Dec 11, 2020

Go to file Add file Code

Template for getting started with automated ML Ops on Azure Machine Learning

Table of repository files and folders including .cloud/.azure, .github/workflows, code, docs/images, .gitignore, CODE_OF_CONDUCT.md, LICENSE, README.md, SECURITY.md.

About

MLOps-Azure Machine Learning

Readme

MIT License

Releases

No releases published Create a new release

Packages

No packages published Publish your first package

Languages




README.md



Search or jump to...

Pull requests Issues Marketplace Explore



rquintino / mlops_azureml Private

Unwatch 1 Star 0 Fork 0

generated from machine-learning-apps/ml-template-azure

Code Issues Pull requests 1 Actions Projects Security Insights Settings

mlops_azureml / .github / workflows / train_deploy.yml

Cancel

Start commit

Edit file Preview changes

Spaces 2 No wrap

```
1 # Actions train a model on Azure Machine Learning
2 name: aml-train-deploy-workflow
3
4 on:
5   push:
6     # branches:
7     #   - master
8     paths:
9       - 'code/*'
10
11 jobs:
12   train:
13     runs-on: ubuntu-latest
14     steps:
15       # Checks-out your repository under $GITHUB_WORKSPACE, so your job can access it
16       - name: Check Out Repository
17         id: checkout_repository
18         uses: actions/checkout@v2
19
20       # Connect or Create the Azure Machine Learning Workspace
21       - name: Connect/Create Azure Machine Learning Workspace
```

Use Control + Space to trigger autocomplete in most situations.

Marketplace Documentation

machine learning

Marketplace / Search results

- Azure Machine Learning Workspace** 15 stars
By Azure Connect to or create an Azure Machine Learning Workspace with this GitHub Action
- Azure Machine Learning Compute Action** 16 stars
By Azure Connect to or create a Compute Target to your Azure Machine Learning Workspace with this GitHub Action
- Azure Machine Learning Run Action** 13 stars
By Azure Submit a run to an Azure Machine

```
19
20 # Connect or Create the Azure Machine Learning Workspace
21 - name: Connect/Create Azure Machine Learning Workspace
22   id: aml_workspace
23   uses: Azure/aml-workspace@v1
24   with:
25     azure_credentials: ${{ secrets.AZURE_CREDENTIALS }}
26
27 # Connect or Create a Compute Target in Azure Machine Learning
28 - name: Connect/Create Azure Machine Learning Compute Target
29   id: aml_compute_training
30   uses: Azure/aml-compute@v1
31   with:
32     azure_credentials: ${{ secrets.AZURE_CREDENTIALS }}
33
34 # Submit a training run to the Azure Machine Learning
35 - name: Submit training run
36   id: aml_run
37   uses: Azure/aml-run@v1
38   with:
39     azure_credentials: ${{ secrets.AZURE_CREDENTIALS }}
40
41 # Register model in Azure Machine Learning model registry
42 - name: Register model
```

master ▾

mlops_azureml / .cloud / .azure / compute.json

Go to file

...



rquintino Update compute.json

Latest commit 8324d31 2 hours ago

[History](#)

1 contributor

8 lines (8 sloc) | 173 Bytes

Raw

Blame




```
1  {
2    "name": "mlopscluster",
3    "compute_type": "amlcluster",
4    "vm_size": "Standard_DS3_v2",
5    "vm_priority": "lowpriority",
6    "min_nodes": 0,
7    "max_nodes": 1
8  }
```



Search or jump to...

Pull requests Issues Marketplace Explore



rquintino / mlops_azureml Private

Unwatch 1

Star 0

Fork 0

generated from machine-learning-apps/ml-template-azure

Code Issues Pull requests 1 Actions Projects Security Insights Settings

master

mlops_azureml / code / train / train.py / Jump to

Go to file



rquintino Initial commit

Latest commit 1c05e79 3 hours ago History

1 contributor

130 lines (103 sloc) 4.59 KB

Raw

Blame



```
1 import os
2 import argparse
3 import itertools
4 import numpy as np
5 import pandas as pd
6 import joblib
7 import matplotlib.pyplot as plt
8
9 from sklearn import datasets
10 from sklearn.svm import SVC
11 from sklearn.metrics import confusion_matrix, precision_score, recall_score, f1_score
12 from sklearn.model_selection import train_test_split
13
14 from azureml.core import Dataset, Run
```

```
x, y = datasets.load_iris(return_X_y=True)
79
80 #dividing X,y into train and test data
81 x_train, x_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=223)
82 data = {'train': {'X': x_train, 'y': y_train},
83         'test': {'X': x_test, 'y': y_test}}
84
85 # train a SVM classifier
86 svm_model = SVC(kernel=args.kernel, C=args.penalty, gamma='auto').fit(data['train']['X'], data['train']['y'])
87 svm_predictions = svm_model.predict(data['test']['X'])
88
89 # accuracy for X_test
90 accuracy = svm_model.score(data['test']['X'], data['test']['y'])
91 print('Accuracy of SVM classifier on test set: {:.2f}'.format(accuracy))
92 run.log('Accuracy', np.float(accuracy))
93
94 # precision for X_test
95 precision = precision_score(svm_predictions, data["test"]["y"], average='weighted')
96 print('Precision of SVM classifier on test set: {:.2f}'.format(precision))
97 run.log('precision', precision)
98
99 # recall for X_test
100 recall = recall_score(svm_predictions, data["test"]["y"], average='weighted')
101 print('Recall of SVM classifier on test set: {:.2f}'.format(recall))
102 run.log('recall', recall)
103
104 # f1-score for X_test
105 f1 = f1_score(svm_predictions, data["test"]["y"], average='weighted')
```

Lets try a new hyper param configuration...

```
127
128 if __name__ == '__main__':
129     args = parse_args()
130     main(args=args)
131
```




Commit changes

change gamma to auto

Add an optional extended description...

- Commit directly to the `master` branch.
- Create a new branch for this commit and start a pull request. [Learn more about pull requests.](#)

 rquintino-gamma-auto

Propose changes

Cancel



Open a pull request

The change you just made was written to a new branch named `rquintino-gamma-auto`. Create a pull request below to propose these changes.



base: master ▾



compare: rquintino-gamma-auto ▾

✓ **Able to merge.** These branches can be automatically merged.



change gamma to auto

Write

Preview



Leave a comment



Attach files by dragging & dropping, selecting or pasting them.



Create pull request

Reviewers



No reviews

Assignees



No one—assign yourself

Labels



None yet

Projects



None yet

Milestone



No milestone

Linked issues



Remember, contributions to this repository should follow its [code of conduct](#).



change gamma to auto #1

rquintino wants to merge 2 commits into master from rquintino-gamma-auto

Add more commits by pushing to the rquintino-gamma-auto branch on rquintino/mlops_azureml.



Some checks haven't completed yet

[Hide all checks](#)

1 in progress check



aml-train-deploy-workflow / train (push) *In progress* — *This check has...*

[Details](#)



This branch has no conflicts with the base branch

Merging can be performed automatically.

Merge pull request



You can also [open this in GitHub Desktop](#)

Projects



None yet

Milestone



No milestone

Linked issues



Successfully merging this pull request may close these issues.

None yet



Write

Preview



Leave a comment

```

1 # Actions train a model on Azure Machine Learning
2 name: aml-train-deploy-workflow
3
4 on:
5   push:
6     # branches:
7     #   - master
8   paths:
9     - 'code/*'
10

```

Attach files by dragging & dropping, selecting or pasting them



Lock conversation

Press **F11** to exit full screen

Summary

Jobs

train

train

Started 33m 3s ago

Search logs



- > Set up job 7s
- > Build Azure/aml-workspace@v1 24s
- > Build Azure/aml-compute@v1 3s
- > Build Azure/aml-run@v1 2s
- > Build Azure/aml-registermodel@v1 21s
- > Build Azure/aml-deploy@v1 3s
- > Check Out Repository 2s
- > Connect/Create Azure Machine Learning Workspace 4s
- > Connect/Create Azure Machine Learning Compute Target 7s
- > Submit training run 31m 10s
- > Register model 9s
- > Deploy model 31s
- > Post Check Out Repository



+ New

Home

Author

Notebooks

Automated ML

Designer

Assets

Datasets

Experiments

Pipelines

Models

Endpoints

Manage

Compute

Datastores

Data Labeling

rq-azureml > Experiments

Experiments

All experiments

All runs



Refresh



Archive experiment



View archived experiments



Search

Experiment	Latest run	Last submitted ↓	Created	Cr
mlops_azureml-master	3	Feb 18, 2021 9:49 PM	Feb 18, 2021 9:24 PM	Se
BankMarketing	11	Jan 21, 2021 4:07 PM	Jan 21, 2021 3:00 PM	Ru

Compute, Experiment, Model, Deployment in Azure ML







+ New

Home

Author

 Notebooks Automated ML Designer

Assets


 Datasets Experiments Pipelines Models Endpoints


Manage

 Compute Datastores Data Labeling

rq-azureml > Compute

Compute

Compute instances Compute clusters Inference clusters Attached compute+ New  Refresh  Delete  View quota Search to filter items...

Name	Provisioning state	Virtual machine size	Created on ↓	Idle nodes	Busy nodes
mlopscluster	 Resizing (0 -> 1 node)	STANDARD_DS3_V2	Feb 18, 2021 9:24 PM	0	0










+ New

Home

Author

 Notebooks Automated ML Designer

Assets

 Datasets Experiments Pipelines Models Endpoints

Manage

 Compute Datastores Data Labeling

rq-azureml > Compute

Compute


Compute instances

Compute clusters

Inference clusters

Attached compute

[+ New](#) [Refresh](#) [Delete](#) [View quota](#)[Search to filter items...](#)

Name	Provisioning state	Virtual machine size	Created on ↓	Idle nodes	Busy nodes
mlopscluster	 Succeeded (1 node)	STANDARD_DS3_V2	Feb 18, 2021 9:24 PM	0	1







+ New

Home

Author

 Notebooks Automated ML Designer




Assets

 Datasets Experiments Pipelines Models Endpoints

Manage

 Compute Datastores Data Labeling

rq-azureml > Experiments > mlops_azureml-master > Run 3

Run 3  Running Refresh  Resubmit  Cancel**Details**

Metrics

Images

Child runs

Outputs + logs

Snapshot

Explanations (preview)

Fairness (preview)

Properties**Status** Running**Created**

Feb 18, 2021 9:49 PM

Started

Feb 18, 2021 10:16 PM

Compute target[mlopscluster](#)**Run ID**

mlops_azureml-master_1613684959_56405295

Script name

train.py

Created by

Service Principal

Input datasets



None

Tags 

GITHUB_ACTOR : rquintino GITHUB_REF : refs/heads/master

GITHUB_REPOSITORY : rquintino/mlops_azureml

GITHUB_SHA : 18b9bce9a9969e24d745430ed8ae238573fd3074

Metrics No data**Description**  Click edit icon to add a description

rq-azureml > Experiments > mlops_azureml-master > Run 3

Run 3  Running Refresh  Resubmit  Cancel | Enable log streaming Word wrapDetails Metrics Images Child runs **Outputs + logs** Snapshot Explanations (preview) Fairness (preview)

confusion_matrix_nc ×



confusion_matrix_ur ×

> azureml-logs

> logs

> outputs

confusion_matrix

  confusion_matrix_normalized_1613686734.png ...  confusion_m confusion_matrix_normalized_1613686734.png

+ New

Home

Author

Notebooks

Automated ML

Designer

Assets

Datasets

Experiments

Pipelines

Models

Endpoints

Manage

Compute

Datastores

Data Labeling



+ New

Home

Author

Notebooks

Automated ML

Designer

Assets

Datasets

Experiments

Pipelines

Models

Endpoints



Manage

Compute

Datastores

Data Labeling

rq-azureml > Experiments > mlops_azureml-master > Run 3

Run 3  Completed Refresh  Resubmit  Cancel**Details**

Metrics

Images


Child runs

Outputs + logs

Snapshot

Explanations (preview)

Fairness (preview)

Properties**Status** Completed**Created**

Feb 18, 2021 9:49 PM

Started

Feb 18, 2021 10:16 PM

Duration

2m 42.44s

Compute target[mlopscluster](#)**Run ID**

mlops_azureml-master_1613684959_56405295

Script name

train.py

Created by

Service Principal

Tags 

GITHUB_ACTOR : rquintino GITHUB_REF : refs/heads/master

GITHUB_REPOSITORY : rquintino/mlops_azureml

GITHUB_SHA : 18b9bce9a9969e24d745430ed8ae238573fd3074

Metrics**confusion_matrix_normalized**

Type: Image

confusion_matrix_unnormalized

Type: Image

confusion_matrix

azureml.v2.confusion_matrix

f1-score

0.9333333333333333



+ New

🏠 Home

Author

📖 Notebooks

⚡ Automated ML

🔧 Designer

Assets

📁 Datasets

🧪 Experiments

🔗 Pipelines

📦 Models

🔄 Endpoints

Manage

💻 Compute

📊 Datastores

📝 Data Labeling

rq-azureml > Experiments > mlops_azureml-master > Run 3

Run 3 ✔ Completed

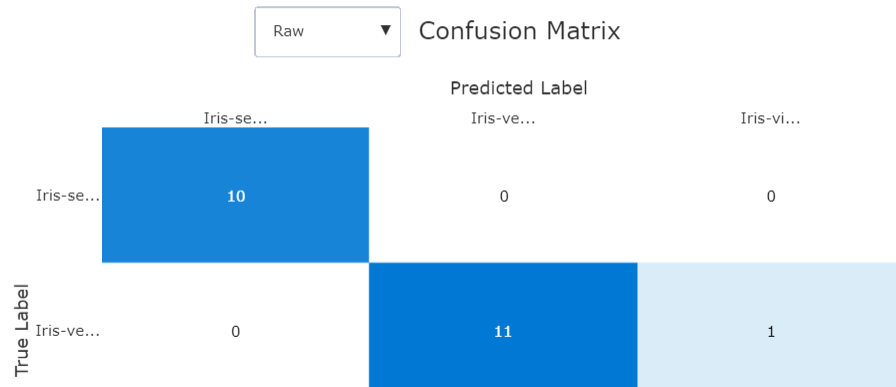

🔄 Refresh ▶ Resubmit ✕ Cancel

Details Metrics Images Child runs Outputs + logs Snapshot Explanations (preview) Fairness (preview)

Select a metric to see a visualization or table of the data.

- Accuracy
- confusion_matrix
- f1-score
- Kernel type
- Penalty
- precision
- recall

View as: Chart Table

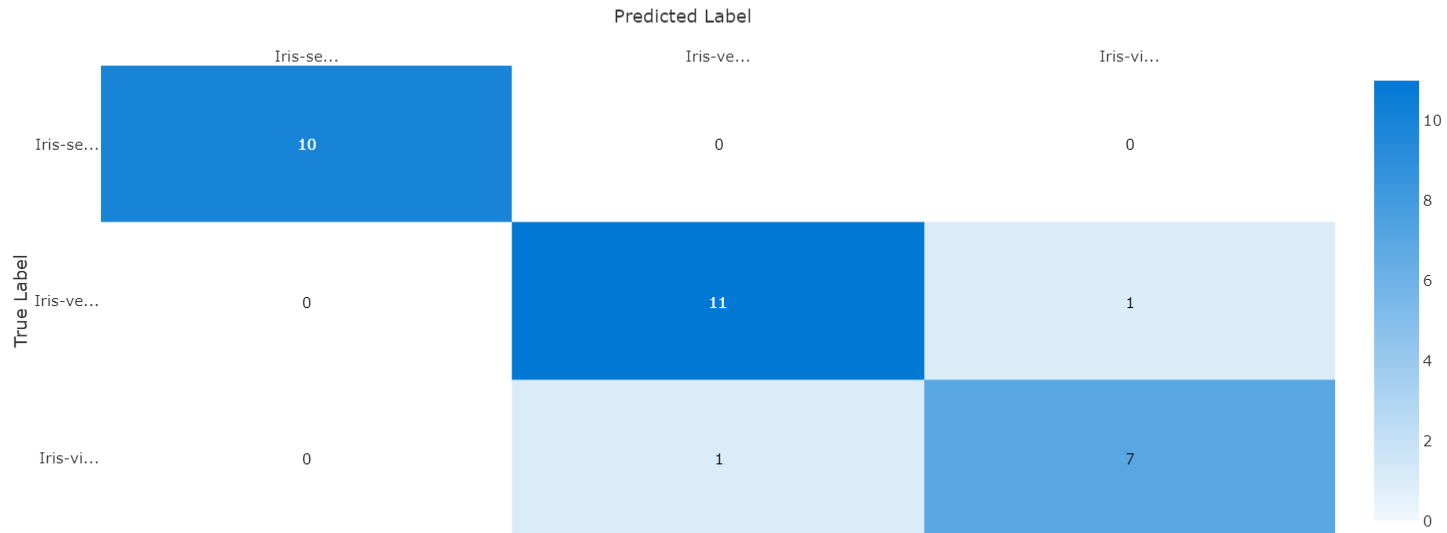




Raw



Confusion Matrix



Summary

Jobs

train

train

succeeded 5 minutes ago in 43m 6s

Search logs




- > ✓ Set up job 7s
- > ✓ Build Azure/aml-workspace@v1 24s
- > ✓ Build Azure/aml-compute@v1 3s
- > ✓ Build Azure/aml-run@v1 2s
- > ✓ Build Azure/aml-registermodel@v1 21s
- > ✓ Build Azure/aml-deploy@v1 3s
- > ✓ Check Out Repository 2s
- > ✓ Connect/Create Azure Machine Learning Workspace 4s
- > ✓ Connect/Create Azure Machine Learning Compute Target 7s
- > ✓ Submit training run 31m 10s
- > ✓ Register model 9s
- > ✓ Deploy model 10m 34s
- > ✓ Post Check Out Repository 0s
- > ✓ Complete job 0s







+ New

Home

Author

 Notebooks Automated ML Designer

Assets



 Datasets Experiments Pipelines **Models** Endpoints

Manage

 Compute Datastores Data Labeling

rq-azureml > Models

Model List

[+ Register model](#)  Delete [▶ Deploy](#)  Refresh

Name	Version	Experiment	Run ID	Created on	Tags
mlops_azureml-master	1	mlops_azureml-master	mlops_azureml-master_161368...	Feb 18, 2021 10:20 PM	



+ New

Home

Author

Notebooks

Automated ML

Designer

Assets

Datasets

Experiments

Pipelines

Models

Endpoints

Manage

Compute

Datastores

Data Labeling

rq-azureml > [Models](#) > mlops_azureml-master:1

mlops_azureml-master:1

Details

Artifacts

Endpoints

Explanations (preview)

Fairness (preview)

Datasets



Refresh



Deploy



Download

Attributes

Version

1

ID

mlops_azureml-master:1

Date registered

2/18/2021, 10:20:33 PM

Framework

Custom


Framework version

--


Experiment name[mlops_azureml-master](#)**Run ID**[mlops_azureml-master_1613684959_56405295](#)**Created by**

6b3d82ad-c37e-4817-9f27-941025e1df47

Tags

 No tags

Properties

 No properties

- ☰
- + New
- 🏠 Home
- Author
- 📓 Notebooks
- ⚡ Automated ML
- 🔗 Designer
- Assets
- 📁 Datasets
- 🧪 Experiments
- 🔗 Pipelines
- 📦 Models**
- 🔄 Endpoints
- Manage
- 💻 Compute
- 📊 Datastores
- 📝 Data Labeling

rq-azureml > Models > mlops_azureml-master:1

mlops_azureml-master:1

Details **Artifacts** Endpoints Explanations (preview) Fairness (preview) Datasets

🔄 Refresh ▶ Deploy ⬇️ Download



📄 model.pkl



📄 model.pkl

ⓘ The file is either binary or uses an unsupported text encoding




Preview not available








+ New

Home

Author

 Notebooks Automated ML Designer

Assets

 Datasets Experiments Pipelines Models Endpoints

Manage



 Compute Datastores Data Labeling

rq-azureml > Endpoints

Endpoints

Real-time endpoints

Pipeline endpoints

 Delete  Edit columns  Refresh  Reset to default view

Showing 1-1 endpoints

Page size:

Name	Description	Created on	Created by	Updated on ↓	Compute type	C
mlops-azureml-master		Feb 18, 2021 10:20 PM	Service Principal	Feb 18, 2021 10:20 PM	ACI	



+ New

Home

Author

Notebooks

Automated ML

Designer

Assets

Datasets

Experiments

Pipelines

Models

Endpoints

Manage

Compute

Datstores

Data Labeling

rq-azureml > Endpoints > mlops-azureml-master

mlops-azureml-master

Details

Test

Consume

Deployment logs

Attributes

Service ID

mlops-azureml-master

Description

--

Deployment state

Transitioning 

Compute type

ACI

Created by

6b3d82ad-c37e-4817-9f27-941025e1df47

Model ID

[mlops_azureml-master:1](#)

Created on


2/18/2021 10:20:44 PM

Last updated on


2/18/2021 10:20:44 PM

Image ID

Tags

 No data

Properties


 No data





+ New

Home


Author

 Notebooks Automated ML Designer

Assets

 Datasets Experiments Pipelines Models Endpoints

Manage

 Compute Datastores Data Labeling

rq-azureml > Endpoints > mlops-azureml-master

mlops-azureml-master

Details

Test

Consume

Deployment logs

Attributes

Service ID

mlops-azureml-master

Description

--

Deployment stateHealthy **Compute type**

ACI

Created by

6b3d82ad-c37e-4817-9f27-941025e1df47

Model ID[mlops_azureml-master:1](#)**Created on**


2/18/2021 10:20:44 PM

Last updated on

2/18/2021 10:20:44 PM

Image ID

Tags

 No data

Properties

hasInferenceSchema

True

hasHttps


False

[+ New](#)[Home](#)[Author](#)[Notebooks](#)[Automated ML](#)[Designer](#)[Assets](#)[Datasets](#)[Experiments](#)[Pipelines](#)[Models](#)[Endpoints](#)[Manage](#)[Compute](#)[Datastores](#)[Data Labeling](#)

rq-azureml > Compute

Compute

[Compute instances](#) **[Compute clusters](#)** [Inference clusters](#) [Attached compute](#)[+ New](#) [Refresh](#) [Delete](#) [View quota](#)[Search to filter items...](#)


Name	Provisioning state	Virtual machine size	Created on ↓	Idle nodes	Busy nodes
mlopscluster	 Succeeded (0 nodes)	STANDARD_DS3_V2	Feb 18, 2021 9:24 PM	0	0



search_models

Experiment ID: 72

Artifact Location: file:///code/data/mlruns/72

Notes 

None

Search Runs:

State

Active


Search

Clear

Showing 29 matching runs

Compare

Delete










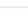
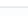
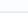
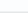






Download CSV 

☰

☰

⚙️

Columns

					Parameters			Metrics >			Tags >			
<input type="checkbox"/>	Date	Run Name	User	Source	Version	FOLD	NOTEBC	SEED	avg_tra	avg_val	avg_val	AILAB_E	AILAB_E	AILAB_E
<input type="checkbox"/>	2021-01-2	MODEL=nasnetalarge BATCH_SIZE=8	rqui...	 ipyke	-	1	1-Tr...	1	1.77...	8.51...	9.00...	202...	doc...	sear...
<input type="checkbox"/>	 2021-0	MODEL=resnet101d BATCH_SIZE=8	rqui...	 ipyke	-	1	1-Tr...	1	4.10...	9.15...	9.66...	202...	doc...	sear...
<input type="checkbox"/>	 2021-0	MODEL=efficientnet_b3a BATCH_SIZE=8	rqui...	 ipyke	-	1	1-Tr...	1	2.13...	8.04...	8.45...	202...	doc...	sear...
<input type="checkbox"/>	 2021-0	MODEL=tresnet_xl_448 BATCH_SIZE=8	rqui...	 ipyke	-	1	1-Tr...	1	-	-	-	202...	doc...	sear...
<input type="checkbox"/>	 2021-0	MODEL=xception71 BATCH_SIZE=8	rqui...	 ipyke	-	1	1-Tr...	1	1.94...	7.92...	8.35...	202...	doc...	sear...
<input type="checkbox"/>	 2021-0	MODEL=xception65 BATCH_SIZE=8	rqui...	 ipyke	-	1	1-Tr...	1	2.01...	8.57...	9.75...	202...	doc...	sear...
<input type="checkbox"/>	 2021-0	MODEL=xception BATCH_SIZE=8	rqui...	 ipyke	-	1	1-Tr...	1	1.96...	7.92...	8.38...	202...	doc...	sear...
<input type="checkbox"/>	 2021-0	MODEL=hrnet_w48 BATCH_SIZE=8	rqui...	 ipyke	-	1	1-Tr...	1	2.46...	8.12...	8.57...	202...	doc...	sear...
<input type="checkbox"/>	 2021-0	MODEL=tf_efficientnet_b8_ap BATCH_SIZE=8	rqui...	 ipyke	-	1	1-Tr...	1	1.76...	8.41...	8.92...	202...	doc...	sear...
<input type="checkbox"/>	 2021-0	MODEL=...	rqui...	 ipyke	-	1	1-Tr...	1	202...	doc...	sear...

Tracking critical for MLOps, but also for Experimenting



Experimentation With GitHub & CML

- Automation for Experimentation
- Ex: Smoke Test Locally (Fast)
- Queue Reproducible in Github
- Using notebooks

master

mlops_azureml / .github / workflows / experiment-train.yml

View runs

Go to file

...



rquintino Update experiment-train.yml

Latest commit b19e525 5 hours ago [History](#)

1 contributor

21 lines (18 sloc) | 526 Bytes

Raw

Blame



```
1 name: experiment-train
2 on: [push]
3
4 jobs:
5   run:
6     runs-on: [ubuntu-latest]
7     container: docker://dvcorg/cml-py3:latest
8     steps:
9       - uses: actions/checkout@v2
10      - name: cml_run
11        env:
12          repo_token: ${ secrets.GITHUB_TOKEN }
13        run: |
14
15          # Your ML workflow goes here
16          pip install -r requirements.txt
17          papermill ./code/experiment/train.ipynb train_output.ipynb
18
19          # Write your CML report
20          cat "Hello" >> report.md
21          cml-send-comment report.md
```

Check Continuous ML (CML) framework in Github

Contents

- 1 Params
- 2 Read data
- 3 Train
- 4 Report

3 Train

```
In [5]: model=RandomForestClassifier()
        executed in 11ms, finished 11:36:38 2021-02-19
```

```
In [6]: cv_scores=cross_val_score(model,X,y,cv=5,scoring="roc_auc")
        cv_scores
        executed in 9.18s, finished 11:36:48 2021-02-19
```

```
Out[6]: array([0.98643411, 0.98013566, 0.99865862, 0.99027498, 0.99865862])
```

4 Report

```
In [7]: mean_scores,std_scores=np.mean(cv_scores),np.std(cv_scores)
        mean_scores,std_scores
        executed in 18ms, finished 11:36:48 2021-02-19
```

```
Out[7]: (0.9908323974856893, 0.007163453627193791)
```

```
In [8]: !echo cv_core_mean: {mean_scores} >> report.md
        !echo cv_core_std: {std_scores} >> report.md
        executed in 212ms, finished 11:36:48 2021-02-19
```

```
In [ ]:
```



Contents

- 1 Params
- 2 Read data
- 3 Train
- 4 Report

TEST=1

executed in 6ms, finished 11:36:38 2021-02-19

In [3]: print(TEST)

executed in 49ms, finished 11:36:38 2021-02-19

1

2 Read data

In [4]: X,y=load_breast_cancer(return_X_y=True)

executed in 94ms, finished 11:36:38 2021-02-19

3 Train

In [5]: model=LogisticRegression()

executed in 1ms, finished 11:36:38 2021-02-19

In [6]: cv_scores=cross_val_score(model,X,y,cv=5,scoring="roc_auc")
cv_scores

executed in 9.18s, finished 11:36:48 2021-02-19

Out[6]: array([[0.98643411, 0.98013566, 0.99865862, 0.99027498, 0.99865862]])

4 Report


Branch for new experiment, change code, queue in GitHub



Search or jump to...

Pull requests Issues Marketplace Explore

125%



rquintino / mlops_azureml Private

generated from machine-learning-apps/ml-template-azure

Unwatch 1 Star 0 Fork 0

<> Code Issues Pull requests 1 Actions Projects Security Insights Settings

Workflows

New workflow

All workflows

aml-train-deploy-workflow

experiment-train

All workflows

Showing runs from all workflows

Filter workflows

6 workflow runs

Event Status Branch Actor

test logistic experiment-train #20: Commit 793afde pushed by rquintino	test-logistic_regression	1 minute ago In progress	...
rf experiment-train #19: Commit 0be9d60 pushed by rquintino	fix_rforest_defaults	3 hours ago 2m 25s	...
fix kernel experiment-train #18: Commit f053740 pushed by rquintino	fix_rforest_defaults	4 hours ago 2m 23s	...
fix kernel experiment-train #14: Commit 10cc9e0 pushed by rquintino	fix_rforest_defaults	4 hours ago 2m 28s	...
Update experiment-train.yml experiment-train #8: Commit f04f092 pushed by rquintino	sklearn-model	5 hours ago 2m 23s	...
Update run_config.yml aml-train-deploy-workflow #6: Commit 18b9bce pushed by rquintino	master	18 hours ago 43m 19s	...

Workflow running...

Summary

Jobs

run

run

succeeded now in 2m 0s

Search logs

```
177 Created wheel for black: filename=black-20.0b1-py3-none-any.whl size=124104 sha256=400c0a0000000000000000000000000000000000000000000000000000000000
178 Stored in directory: /tmp/pip-ephem-wheel-cache-zlkyu1xx/wheels/dd/c8/94/e774504958e4667180d81202did4d5cdc35d6222ae39229871b
179 Successfully built black
180 Installing collected packages: zipp, typing-extensions, six, ipython-genutils, decorator, traitlets, pyrsistent, importlib-metadata, attrs, wcwidth,
tornado, pyzmq, python-dateutil, pyarsing, ptyprocess, parso, jupyter-core, jsonschema, webencodings, webencodings, weenschema, prompt-toolkit, pickleshare,
pexpect, packaging, nest-asyncio, nbformat, MarkupSafe, jupyter-client, jedi, backcall, async-generator, urllib3, typed-ast, tomli, textwrap3, testpath,
regex, pathspec, pandocfilters, numpy, nbclient, mpy-extensions, mistune, jupyterlab-pygments, Jinja2, ipython, idna, entrypoints, defusedxml, dataclasses,
click, chardet, cffi, certifi, bleach, appdirs, tqdm, ThreadPoolCtl, terminado, tenacity, Send2Trash, scipy, requests, pyyaml, prometheus-client, nbconvert,
joblib, ipykernel, black, argon2-cffi, ansiwrap, scikit-learn, papermill, notebook
181 Running setup.py install for pyrsistent: started
182 Running setup.py install for pyrsistent: finished with status 'done'
183 Running setup.py install for pandocfilters: started
184 Running setup.py install for pandocfilters: finished with status 'done'
185 Successfully installed MarkupSafe-1.1.1 Send2Trash-1.5.0 ansiwrap-0.8.4 appdirs-1.4.4 argon2-cffi-20.1.0 async-generator-1.10 attrs-20.3.0 backcall-0.2.0
black-20.0b1 bleach-3.3.0 certifi-2020.12.5 cffi-1.14.5 chardet-4.0.0 click-7.1.2 dataclasses-0.8 decorator-4.4.2 defusedxml-0.6.0 entrypoints-0.3 idna-2.10
importlib-metadata-3.4.0 ipykernel-5.4.3 ipython-7.16.1 ipython-genutils-0.2.0 jedi-0.18.0 Jinja2-2.11.3 joblib-1.0.1 jsonschema-3.2.0 jupyter-client-6.1.11
jupyter-core-4.7.1 jupyterlab-pygments-0.11.2 mistune-0.8.4 mpy-extensions-0.4.3 nbclient-0.5.2 nbconvert-6.0.7 nbformat-5.1.2 nest-asyncio-1.5.1 notebook-
6.2.0 numpy-1.19.5 packaging-20.9 pandocfilters-1.4.3 papermill-2.3.2 parso-0.8.1 pathspec-0.8.1 pexpect-4.8.0 pickleshare-0.7.5 prometheus-client-0.9.0
prompt-toolkit-3.0.16 ptyprocess-0.7.0 pycparser-2.20 pygments-2.8.0 pyarsing-2.4.7 pyrsistent-0.17.3 python-dateutil-2.8.1 pyyaml-5.4.1 pyzmq-22.0.3
regex-2020.11.13 requests-2.25.1 scikit-learn-0.24.1 scipy-1.5.4 six-1.15.0 tenacity-6.3.1 terminado-0.9.2 testpath-0.4.4 textwrap3-0.9.2 ThreadPoolCtl-
2.1.0 tomli-0.10.2 tornado-6.1 tqdm-4.57.0 traitlets-4.3.3 typed-ast-1.4.2 typing-extensions-3.7.4.3 urllib3-1.26.3 wcwidth-0.2.5 webencodings-0.5.1 zipp-
0.6.0
```

```
186 Input Notebook: code/experiment/train.ipynb
187 Output Notebook: train_output.ipynb
188
189 Executing: 0% | 0/13 [00:00<, 2cell/s]Executing notebook with kernel: python3
190
191 Executing: 8% | 1/13 [00:01<00:15, 1.29s/cell]
192 Executing: 54% | 7/13 [00:01<00:00, 6.66cell/s]
193 Executing: 85% | 11/13 [00:01<00:00, 9.64cell/s]
194 Executing: 100% | 13/13 [00:03<00:00, 4.09cell/s]
```

- Post Run actions/checkout@v2 0s
- Stop containers 6s
- Complete job 0s

Notebook execution using papermill on Github agent

```
15 "from sklearn.model_selection import train_test_split\n",
16 "from sklearn.datasets import load_breast_cancer\n",
17 "from sklearn.ensemble import RandomForestClassifier\n",
18 "import numpy as np"
19 ]
20 },
110 },
111 "outputs": [],
112 "source": [
113 - "model=RandomForestClassifier()"
114 ]
115 },
116 {
117 + "from sklearn.model_selection import train_test_split\n",
118 + "from sklearn.datasets import load_breast_cancer\n",
119 + "from sklearn.ensemble import RandomForestClassifier\n",
120 + "from sklearn.linear_model import LogisticRegression\n",
121 + "import numpy as np"
122 ]
123 },
124 "outputs": [],
125 "source": [
126 + "model=LogisticRegression()"
127 ]
128 },
129 }
```

1 comment on commit 793afde

github-actions bot commented on 793afde 6 minutes ago

Train Report

cv_core_mean: 0.9910938715843727
cv_core_std: 0.006007611645631809

Train Finished

Re: [rquintino/mlops_azureml] test logistic (793afde) Inbox x

github-actions[bot] <notifications@github.com> [Unsubscribe](#)
to rquintino/mlops_azureml, Subscribed

Train Report

cv_core_mean: 0.9910938715843727
cv_core_std: 0.006007611645631809

Train Finished

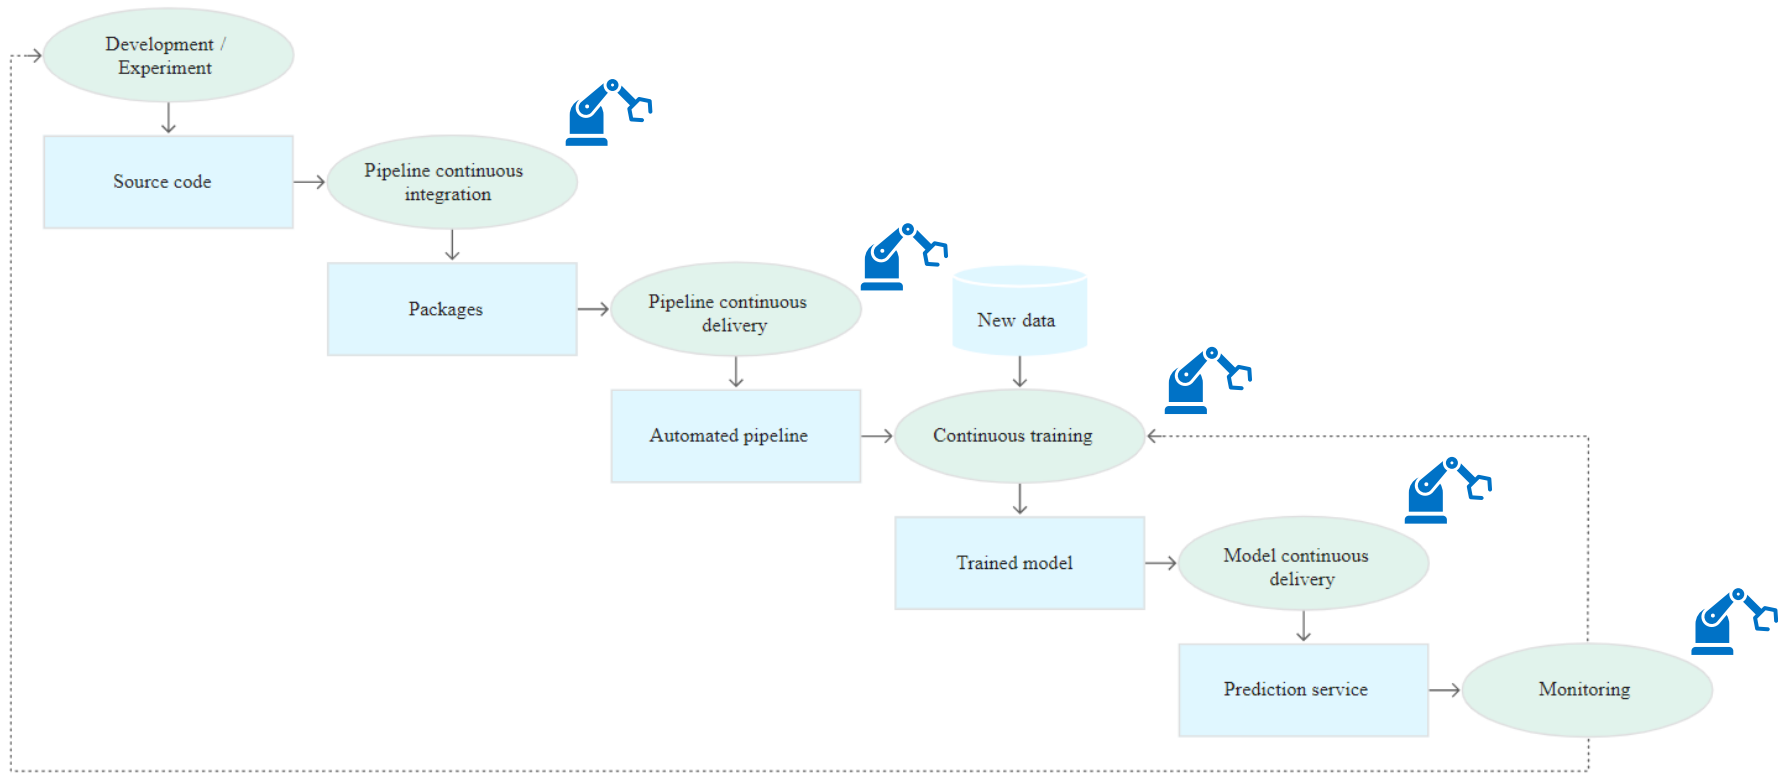
[CML watermark](#)

—
You are receiving this because you are subscribed to this thread.
Reply to this email directly, [view it on GitHub](#), or [unsubscribe](#).

Train performance results on commits (also PRs)

MLOps Recap

- From Models to Pipelines
- Execution of Pipelines produce models
- Train for Deployment is fully automated/reproducible
- Pipelines enable auto model retraining
- Automation, Automation, Automation...
- Monitoring for Data Drift, Concept Drift, Model Drift
- Tracking, Metadata & Artifact Stores



Adapted from: <https://cloud.google.com/solutions/machine-learning/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning>

Mature MLOps Example (Level 1)

Run 3 Queued

Refresh Resubmit Cancel | Enable log streaming Word wrap

Details Metrics Images Child runs **Outputs + logs** Snapshot Explanations (preview) Fairness (preview)



20_image_build_log. x

azureml-logs

20_image_build_log.txt

```
670 latest: digest: sha256:7df4cc95e19514ea3b95cb16faf838c26d4a32e786553383ea5c619d975!
671 2021/02/18 22:12:24 Successfully pushed image: 05cf3362d2b44e3a97168a21e22f74e6.azurecr.io
672 2021/02/18 22:12:24 Step ID: acb_step_0 marked as successful (elapsed time in :
673 2021/02/18 22:12:24 Populating digests for step ID: acb_step_0...
674 2021/02/18 22:12:25 Successfully populated digests for step ID: acb_step_0
675 2021/02/18 22:12:25 Step ID: acb_step_1 marked as successful (elapsed time in :
676 2021/02/18 22:12:25 The following dependencies were found:
677 2021/02/18 22:12:25
678 - image:
679   registry: 05cf3362d2b44e3a97168a21e22f74e6.azurecr.io
680   repository: azureml/azureml_b40693e99d8d288ddefab00c6595c11c
681   tag: latest
682   digest: sha256:7df4cc95e19514ea3b95cb16faf838c26d4a32e786553383ea5c619d975!
683 runtime-dependency:
684   registry: mcr.microsoft.com
685   repository: azureml/base
686   tag: intelmpi2018.3-ubuntu16.04
687   digest: sha256:a1b514f3ba884b9a7695cbbba5638933dda222e8ce3e8c81e8cdf861679:
688   git: {}
689
690 Run ID: cb3 was successful after 22m55s
691
```

Not best fit for interative/model development

Final notes

- MLOps doesn't guarantee ML success
- MLOps should improve feedback and learning
(not block DataScientists from business/data understanding)
- Dont reinvent the wheel...
- Still a young field
Dont believe everything you read in docs (and ppt...)
- When/If to introduce MLOps
(it depends...)



References

MLOps



Press F11 to exit full screen

Machine Learning Ops

A collection of resources on how to facilitate Machine Learning Ops with GitHub.



Learn how to use GitHub for automation, collaboration and reproducibility in your machine learning workflows.

<https://mlops.githubapp.com/>

DSPT #84



Catarina Silva

- Senior Research Engineer @ Unbabel
- Mother of GPUs, freer of CPUs and rightful queen of NMT
- Star wars fan, Engineer, Developer — in that order.



2:58 / 1:17:14



DSPT#84 Webinar- The rise of MLOps

494 views • Streamed live on Feb 3, 2021



31



0



SHARE



SAVE




LICENSE	initial commit	12 months ago
README.md	Contrib/Fix typos (#395) (#396)	5 days ago
bitbucket-pipelines.yml	Bb pr comments (#347)	3 months ago
package-lock.json	Remove new line coerce in private ssh (#408)	18 hours ago
package.json	Remove new line coerce in private ssh (#408)	18 hours ago
prettier.config.js	Pretifier update (#239)	6 months ago

Contributors 17



+ 6 contributors

Languages



README.md



What is CML? Continuous Machine Learning (CML) is an open-source library for implementing continuous integration & delivery (CI/CD) in machine learning projects. Use it to automate parts of your development workflow, including model training and evaluation, comparing ML experiments across your project history, and monitoring changing datasets.



andronovhopf commented on c5eebf3 21 hours ago

Author Member

Style transfer

Workspace vs. Master



Training metrics

Path	Param	Old	New
params.yaml	iter	1	30

GPU info

Also


- <https://azure.microsoft.com/en-us/services/machine-learning/mlops/>
- <https://cloud.google.com/solutions/machine-learning/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning>
- <https://github.com/microsoft/MLOps>
- <https://databricks.com/blog/2020/10/13/using-mlops-with-mlflow-and-azure.html>
- <https://github.com/Azure/MachineLearningNotebooks>

Thank you!

Q&A

 ruí.quintino@devscope.net

 [Linkedin.com/ in/rquintino](https://www.linkedin.com/in/rquintino)



devscope



Rua Passos Manuel Nº 223 – 4º Andar
4000-385 Porto

Av. Sidónio Pais, Nº 2 – 3º Andar
1050-214 5 Lisboa



T. +351 223 751 350/51
F. +351 223 751 352



info@devscope.net
www.devscope.net

Helping customers succeed



