

THE DEVELOPER'S CONFERENCE

Trilha – Microsoft

MSc. José Ahirton Lopes

Chief Data Officer @ Lambda3

Professor de MBA @ FIAP

MVP AI @ Microsoft

PhD Candidate @ Mackenzie

Produzindo Machine Learning em Escala: WTF Is MLOps?

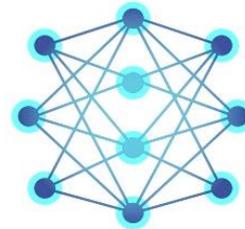


THE
DEVELOPER'S
CONFERENCE

MSc. José Ahirton Lopes



LAMBDA3



AI Brasil

Qual a Escala?



THE
DEVELOPER'S
CONFERENCE

**180
million**

**Monthly active
Office 365 users
using AI**

**18
Billion**

**Questions Asked
of Cortana**

**6.5
Trillion**

**Number of Signals
Analyzed to Block
Emerging Threats
DAILY**

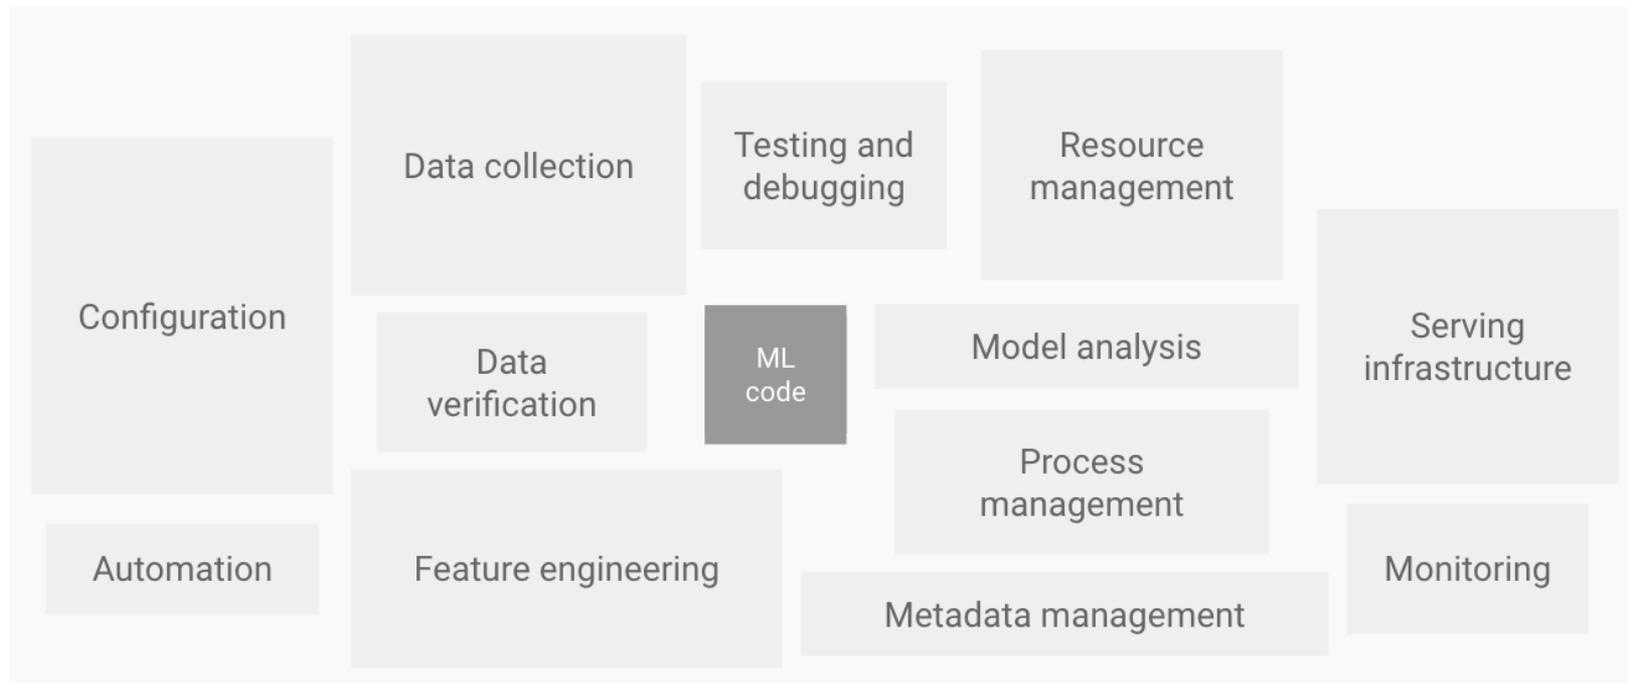
ML é muito mais que nossos
modelos!



THE
DEVELOPER'S
CONFERENCE

Building a model

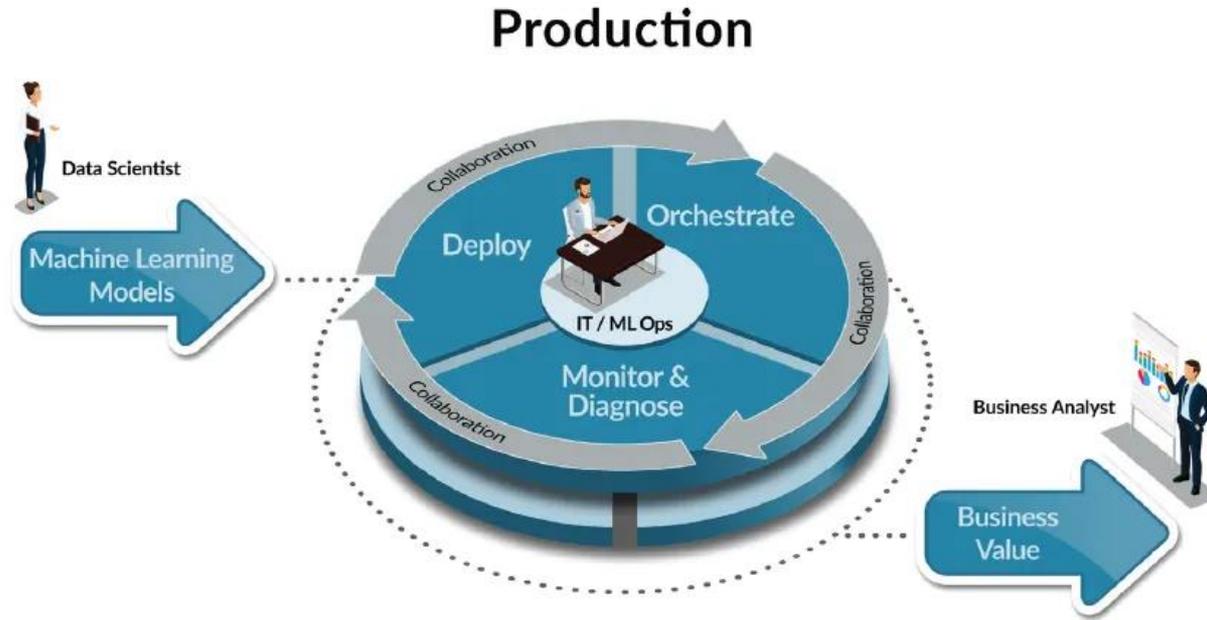
ML é muito mais que nossos modelos!



ML só vira *insight* quando em produção



THE
DEVELOPER'S
CONFERENCE



The ML Management Lifecycle enabled by ParallelM

ML só vira *insight* quando em produção



THE
DEVELOPER'S
CONFERENCE



ginablaber

@ginablaber

Follow



The story of enterprise Machine Learning: “It took me 3 weeks to develop the model. It’s been >11 months, and it’s still not deployed.”
[@DineshNirmalIBM](#) [#StrataData](#) [#strataconf](#)

10:19 AM - 7 Mar 2018

7 Retweets 19 Likes

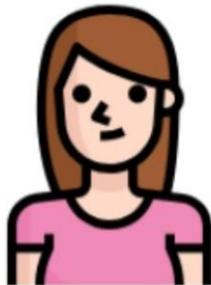


↻ 7



19

Precisamos de ferramentas e métodos!



Data Scientist

- Quick iteration
- Frameworks they understand
- Best of breed tools
- No management headaches
- Unlimited scale



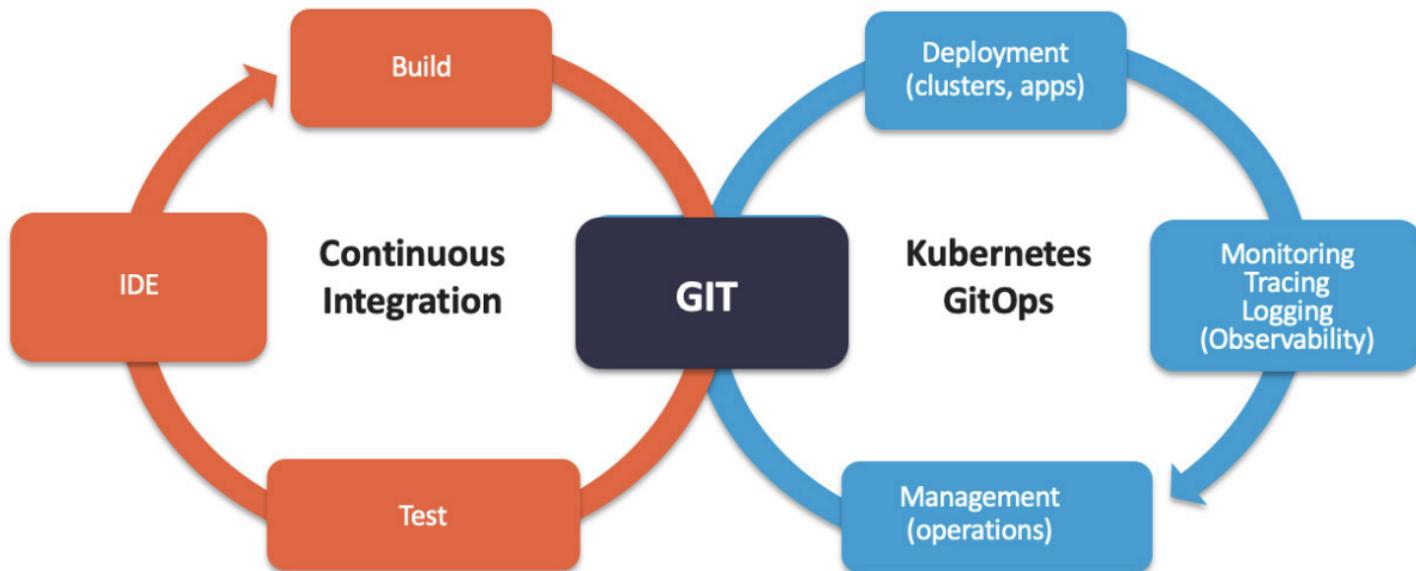
SRE/ML Engineers

- Reuse of tooling and platforms
- Corporate compliance
- Observability
- Uptime

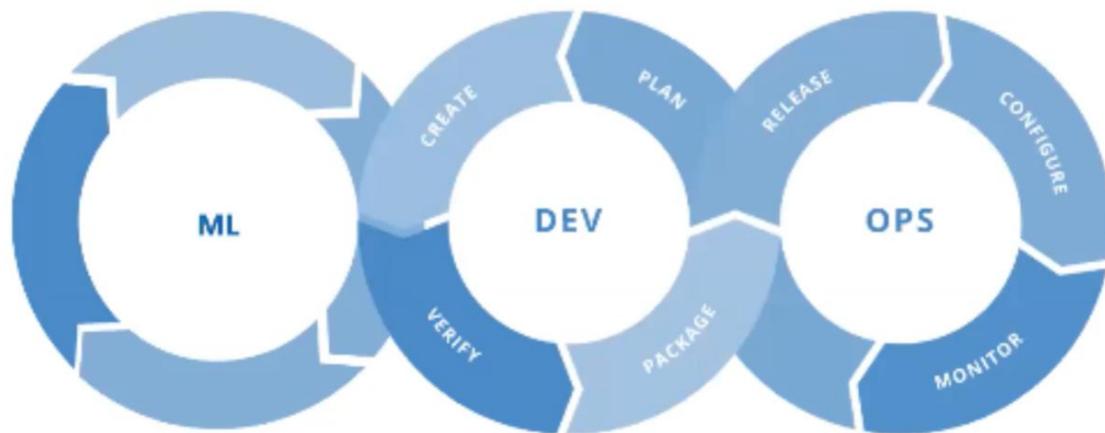
GitOps



THE
DEVELOPER'S
CONFERENCE



MLOps = ML + DEV + OPS



Experiment

Data Acquisition
Business Understanding
Initial Modeling

Develop

Modeling + Testing
Continuous Integration
Continuous Deployment

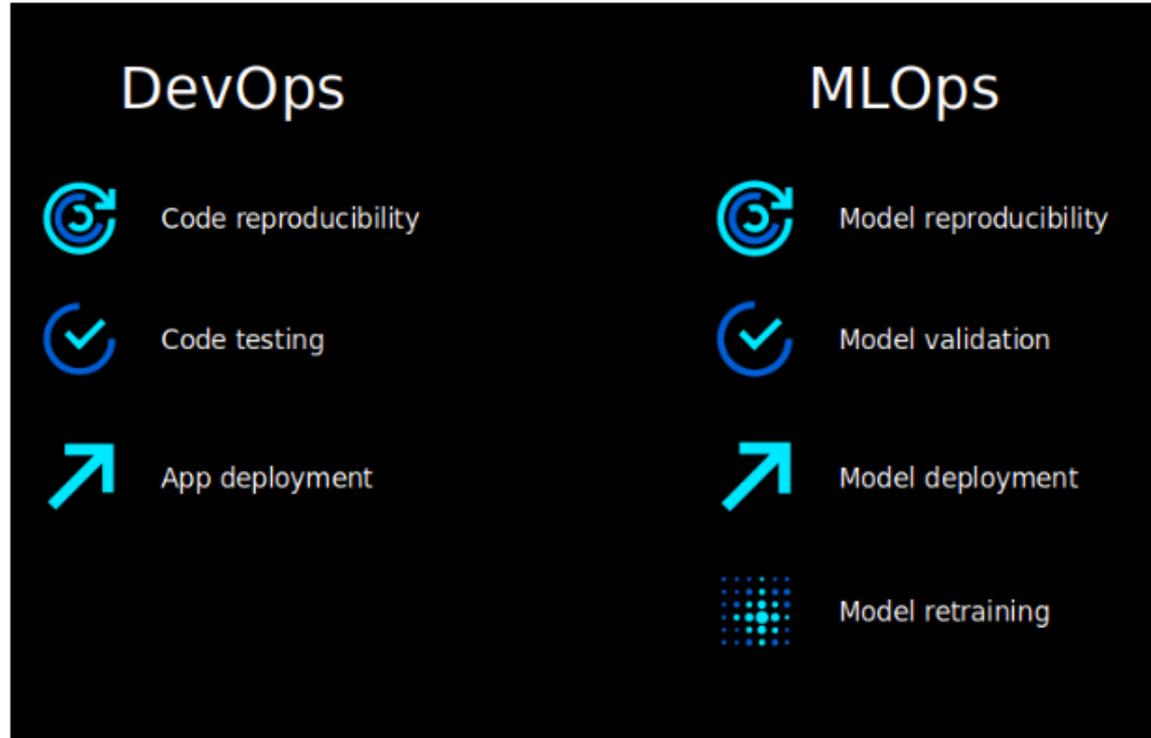
Operate

Continuous Delivery
Data Feedback Loop
System + Model Monitoring

DevOps e MLOps?



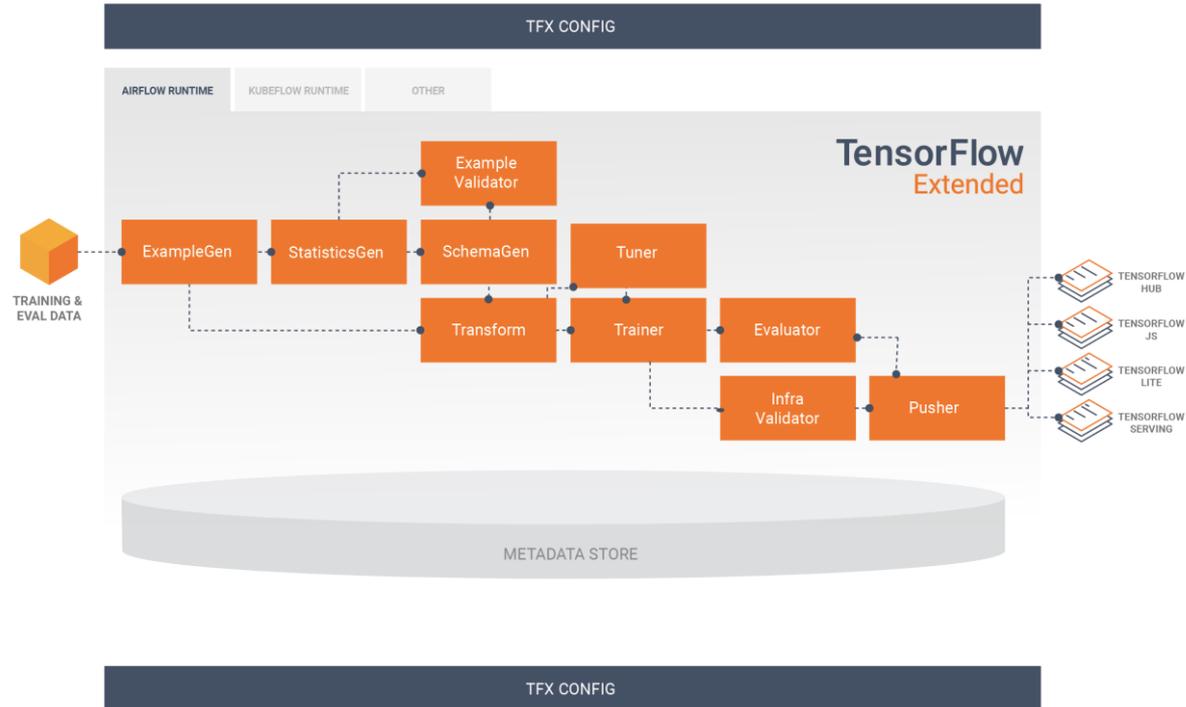
THE
DEVELOPER'S
CONFERENCE



Plataformas de MLOps – TensorFlow Extended



THE
DEVELOPER'S
CONFERENCE



Plataformas de MLOps – FBLearner Flow



THE
DEVELOPER'S
CONFERENCE

FBLearner Workflow Library Models/Features Help Search for people, tasks, tools...

FBLearner Search by workflow, tag, owner, name, or ID Advanced search

Launch New Run Compare

My Runs My Test Runs My Recurring Runs All Runs Custom +

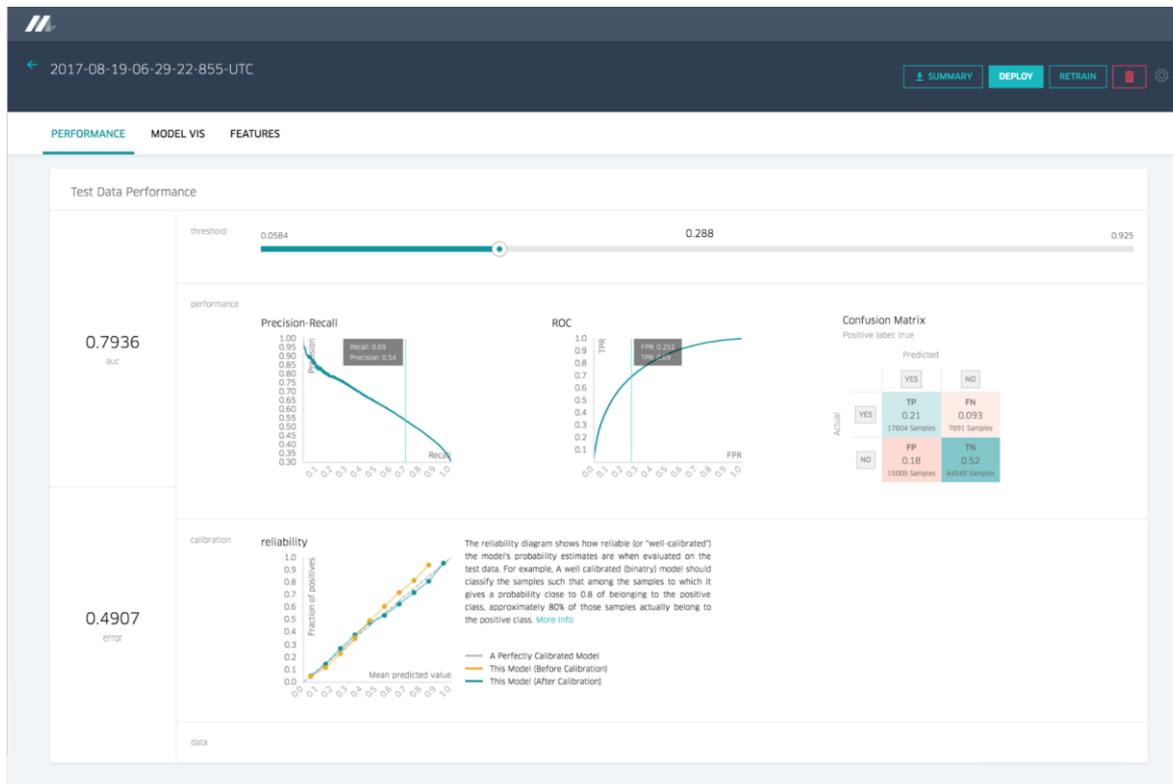
ID	Owner	Workflow	Name	Progress	Start Time	Tags	Log Loss	AUC
1047165	Mahaveer Jain	Parameter Sweep Example	-	<div style="width: 100%;"></div>	9/9, 9:06pm	london-demo	-	-
1047298	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.35	<div style="width: 100%;"></div>	9/9, 9:19pm	-	0.00105	0.95524
1047297	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.25	<div style="width: 100%;"></div>	9/9, 9:19pm	-	0.00107	0.95776
1047296	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.3	<div style="width: 100%;"></div>	9/9, 9:19pm	-	0.00104	0.95719
1047295	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.1	<div style="width: 100%;"></div>	9/9, 9:19pm	-	0.00122	0.95871
1047294	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.2	<div style="width: 100%;"></div>	9/9, 9:19pm	-	0.00109	0.95796
1047293	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.15	<div style="width: 100%;"></div>	9/9, 9:19pm	-	0.00115	0.95887
1047292	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.4	<div style="width: 100%;"></div>	9/9, 9:19pm	-	0.00106	0.95355
1047291	Mahaveer Jain	Gradient Boosted Decision Tree Training	Learning Rate: 0.45	<div style="width: 100%;"></div>	9/9, 9:19pm	-	0.00110	0.95293
1037778	Jason Briceno	Parameter Sweep Example	-	<div style="width: 100%;"></div>	9/8, 2:30pm	-	-	-
950428	Li Zhang	Parameter Sweep Example	-	<div style="width: 100%;"></div>	8/21, 2:40pm	-	-	-
900673	Jiawei Chen	Parameter Sweep Example	-	<div style="width: 100%;"></div>	8/8, 9:11pm	-	-	-
832281	Giri Rajaram	Parameter Sweep Example	-	<div style="width: 100%;"></div>	7/24, 12:56pm	-	-	-
832027	Giri Rajaram	Parameter Sweep Example	-	<div style="width: 100%;"></div>	7/24, 12:34pm	-	-	-

Displaying results 1 - 8 out of 8 matches.

Plataformas de MLOps – UBER Michelangelo



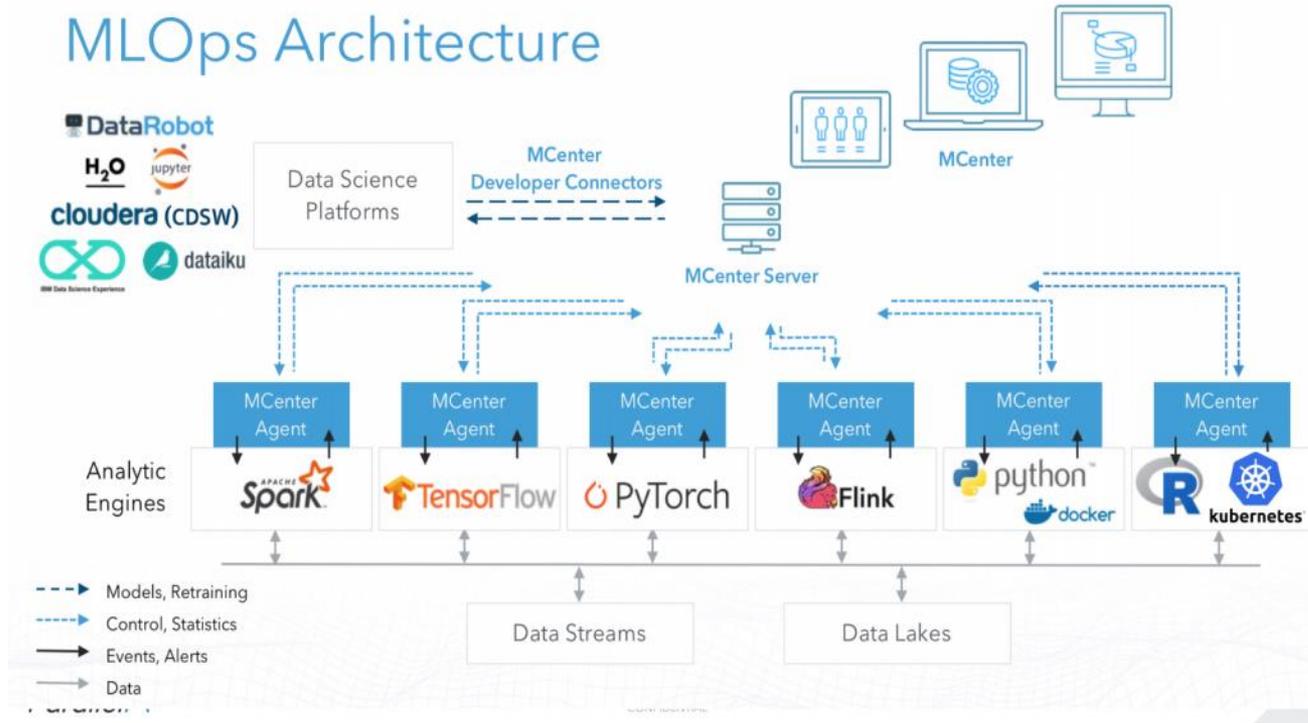
THE
DEVELOPER'S
CONFERENCE



Plataformas de MLOps – Por que não construir?



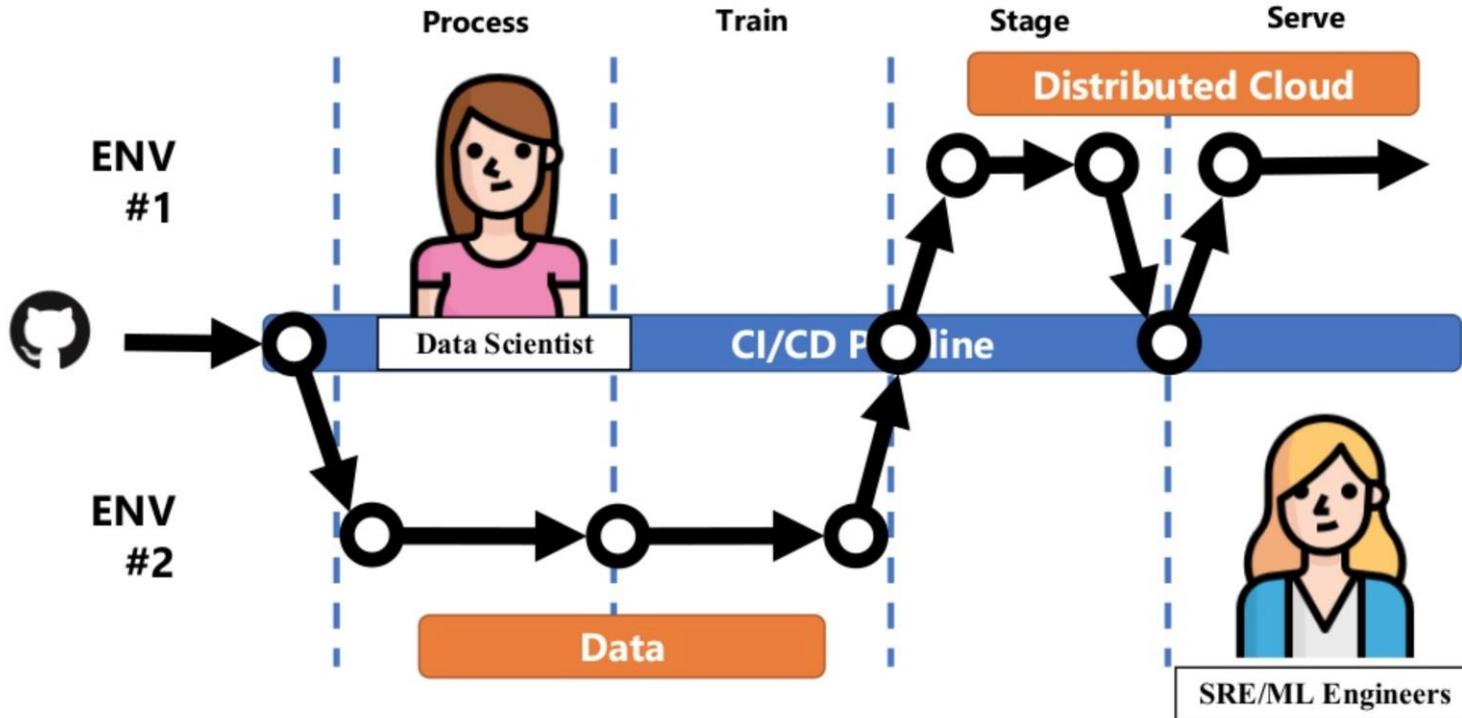
MLOps Architecture



Ambientes Multi-Cloud



THE
DEVELOPER'S
CONFERENCE



Pipelines no Azure DevOps



THE
DEVELOPER'S
CONFERENCE

Cloud-hosted pipelines for Linux, Windows and macOS.



Any language, any platform, any cloud

Build, test, and deploy Node.js, Python, Java, PHP, Ruby, C/C++, .NET, Android, and iOS apps. Run in parallel on Linux, macOS, and Windows. Deploy to Azure, AWS, GCP or on-premises



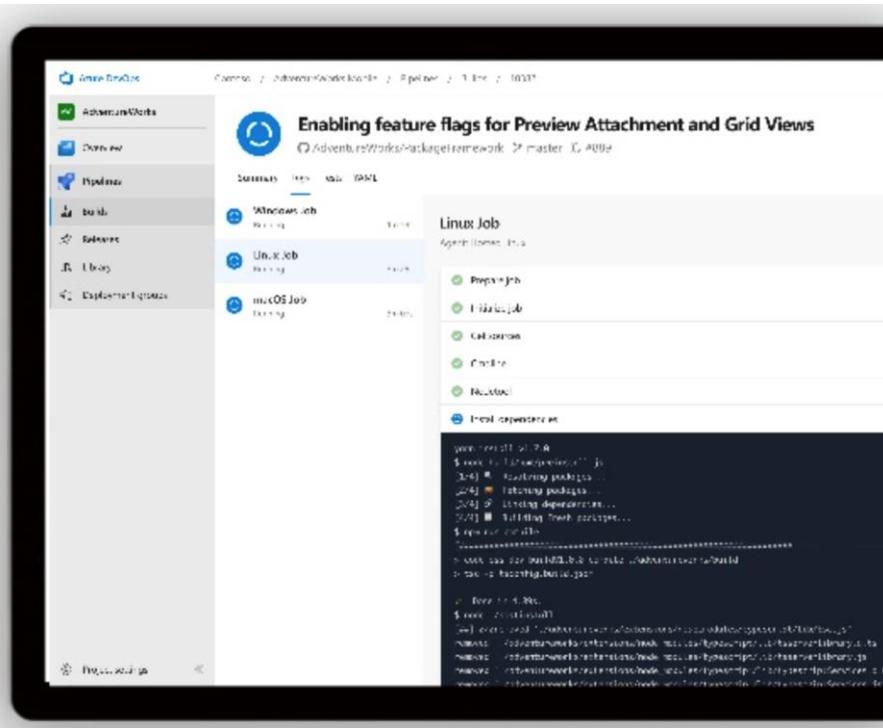
Extensible

Explore and implement a wide range of community-built build, test, and deployment tasks, along with hundreds of extensions from Slack to SonarCloud. Support for YAML, reporting and more



Containers and Kubernetes

Easily build and push images to container registries like Docker Hub and Azure Container Registry. Deploy containers to individual hosts or Kubernetes.



Azure DevOps + Azure ML

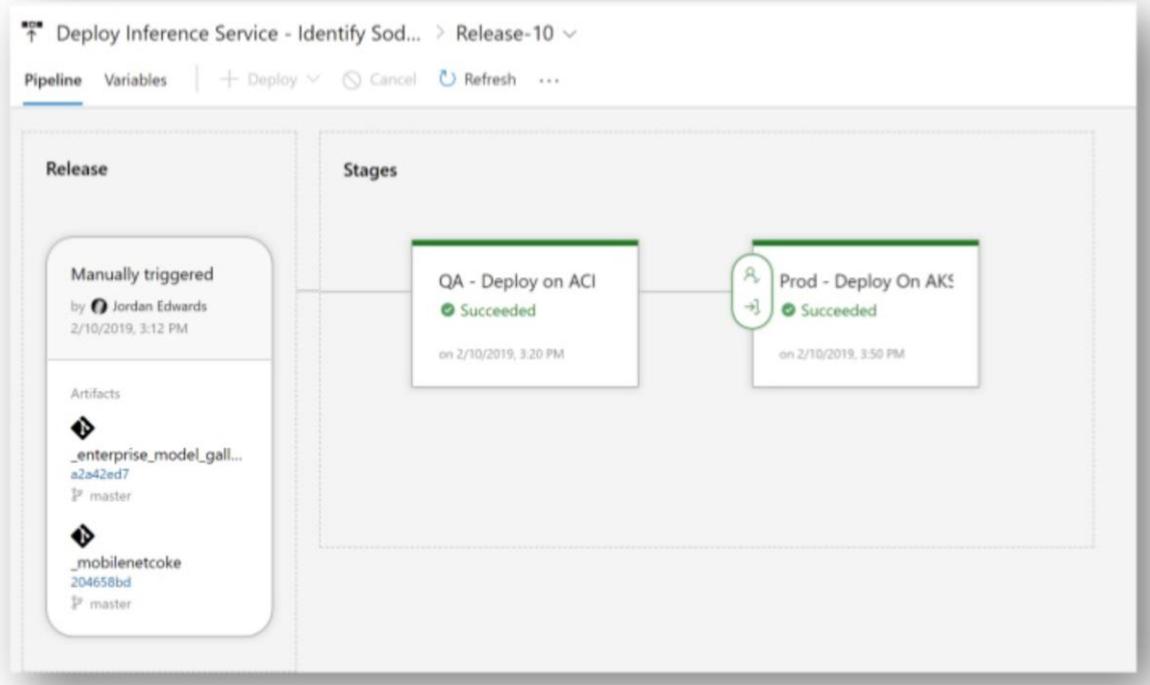


Trigger release pipelines
on model registration

Use **Azure DevOps +
Azure ML CLI** to manage
E2E release flow

Leverage Azure DevOps
approvals for controlling
application rollout

**Announcing: The AML
Extension for MLOps**



Deploy Inference Service - Identify Sod... > Release-10

Pipeline Variables | + Deploy | Cancel | Refresh | ...

Release

Manually triggered
by Jordan Edwards
2/10/2019, 3:12 PM

Artifacts

- _enterprise_model_gall...
a2e42ed7
master
- _mobilenetcoke
204658bd
master

Stages

- QA - Deploy on ACI
Succeeded
on 2/10/2019, 3:20 PM
- Prod - Deploy On AKS
Succeeded
on 2/10/2019, 3:50 PM

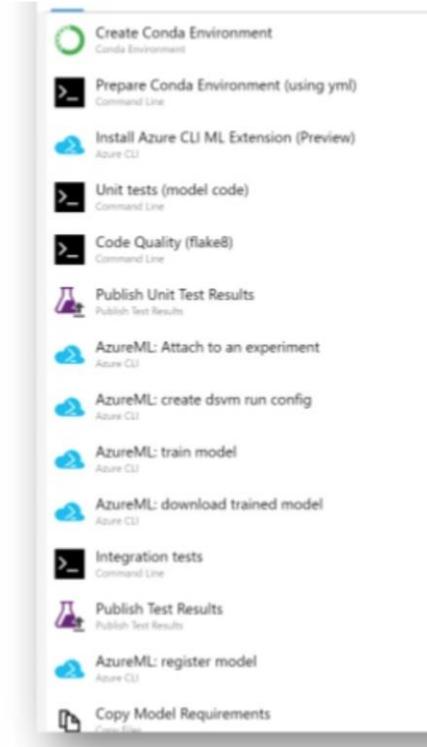
Treinamento de Modelos



THE
DEVELOPER'S
CONFERENCE

CI pipeline captures:

1. Create sandbox
2. Run unit tests and code quality checks
3. Attach to compute
4. Run training pipeline
5. Evaluate model
6. Register model



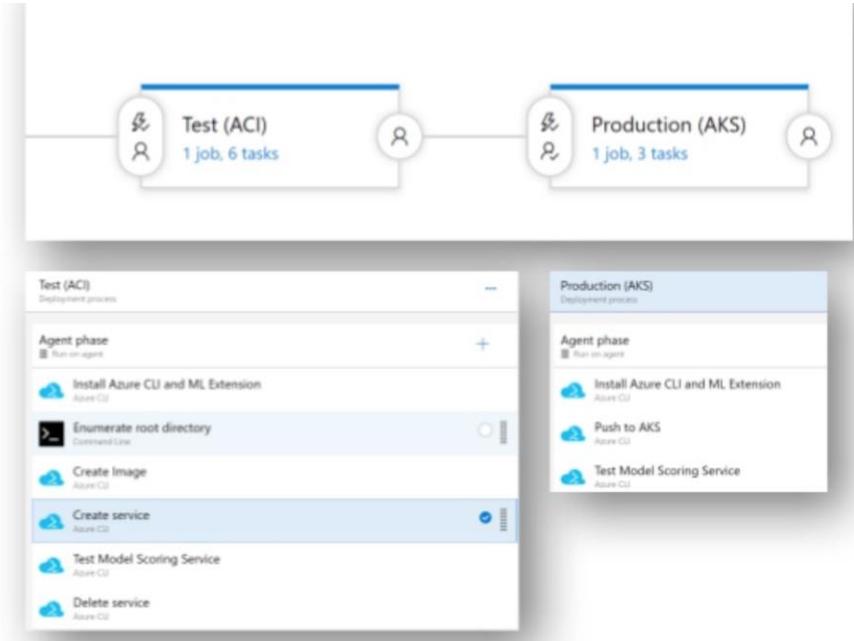
Deploy Automatizado



CD pipeline captures:

1. Package model into container image
2. Validate and profile model
3. Deploy model to DevTest (ACI)
4. If all is well, proceed to rollout to AKS

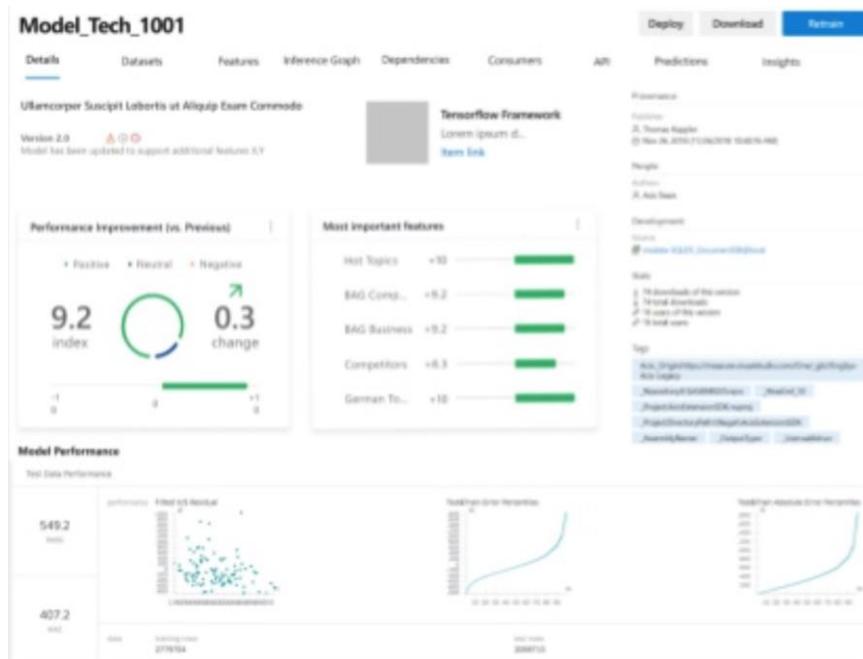
Everything is done via the CLI



Versionamento e Armazenamento



THE
DEVELOPER'S
CONFERENCE



Validação de Modelos



THE
DEVELOPER'S
CONFERENCE

```
In [24]: !curl --request POST --data-binary "@test-image\pepsi.jpg" http://52.157.23.225:80/score
```

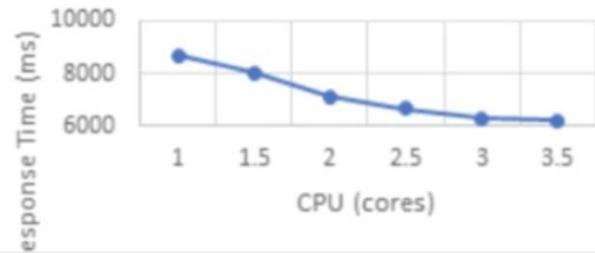
```
"[{"pepsi": "1.0"}, {"ice": "2.5673091e-10"}, {"coke": "2.759612e-13"}]"
```

% Total	% Received	% Xferd	Average Speed	Time	Time	Time	Current
			Dload Upload	Total	Spent	Left	Speed
0	0	0	0	0	--:--:--	--:--:--	0
0	185k	0	0	0	--:--:--	--:--:--	0
100	185k	0	0	100	185k	0	86045
100	185k	100	84	100	185k	29	66302
100	185k	100	84	100	185k	29	66302

Response Time vs Memory



Response Time vs CPU



Deployment de Modelos



THE
DEVELOPER'S
CONFERENCE

The screenshot displays the Azure DevOps interface for a release named 'Example Scikit Release' (Release-11). The pipeline consists of two stages: 'ACI Test Deployment' and 'AKS PROD Deployment', both of which have succeeded. The 'Release' section on the left shows it was manually triggered by Joel Roggevan on 6/26/2018 at 5:57 PM. Below this, an artifact named 'models' is listed with a version of '0704e4d0' and a 'master' branch. A detailed view of the 'AKS PROD Deployment' stage is shown in the foreground, listing various service properties.

Property	Value
Service ID	gpt9g9ev0alv0m0v0b0v0n02012101000000
Creation date	5/1/2018, 12:42:28 PM
Last updated	5/1/2018, 12:42:28 PM
Status	Succeeded
Environment	gpt9g9ev0alv0m0v0b0v0n02012101000000
URL	http://10.0.0.1:8080/
Primary key	gpt9g9ev0alv0m0v0b0v0n02012101000000
Secondary key	gpt9g9ev0alv0m0v0b0v0n02012101000000
Event hub	None
Storage	True
CPU reserve capacity	0.1
Memory reserve capacity	5000
Scale settings	None
Scale type	Static

MLOps te leva onde você quer chegar



THE
DEVELOPER'S
CONFERENCE

- Empoderamento ponta a ponta para suas equipes de CD usarem as melhores práticas;
- Entrega contínua de real valor aos usuários;
- Promove auditabilidade, compliance junto a órgãos regulatórios (alguém falou LGPD?) e consistência!

O que resolve (Moo!)?



- **M**eu modelo realmente funciona?
- **O** que quem está na ponta está vendo?
- **O** meu modelo ainda é bom?

Por que se preocupar com isso?



How We Analyzed the COMPAS Recidivism Algorithm

by Jeff Larson, Surya Mattu, Lauren Kirchner and Julia Angwin

May 23, 2016

[← Read the story](#)

Across the nation, judges, probation and parole officers are increasingly using algorithms to assess a criminal defendant's likelihood of becoming a recidivist – a term used to describe criminals who re-offend. There are dozens of these risk assessment algorithms in use. Many states have built their own assessments, and several academics have written tools. There are also two leading nationwide tools offered by commercial vendors.

Justiça dá 30 dias para que Metrô de SP esclareça projeto de câmeras de reconhecimento facial

Metrô deverá explicar a falta de debate público prévio e apresentar os estudos que utilizou para medir os benefícios e riscos do sistema. Em outubro foi escolhido consórcio para implantar sistema com proposta de R\$ 58,6 milhões.

Por G1 SP — São Paulo

12/02/2020 20h37 · Atualizado há um ano



O que vem por aí?





THE
DEVELOPER'S
CONFERENCE

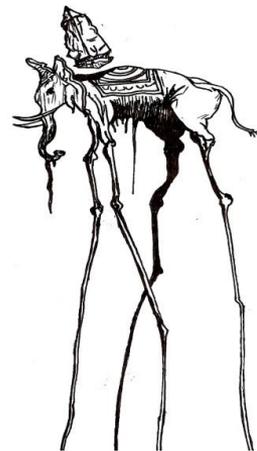
Como aprender MLOps com a Azure?



THE
DEVELOPER'S
CONFERENCE

*“É preciso provocar
sistematicamente confusão. Isso
promove a criatividade. Tudo aquilo
que é contraditório gera vida ”*

Salvador Dalí



Meus Contatos



THE
DEVELOPER'S
CONFERENCE



LinkedIn



GitHub



ahirtonlopes@gmail.com

Contatos – AI Brasil



Linked in



meetup



WhatsApp



facebook

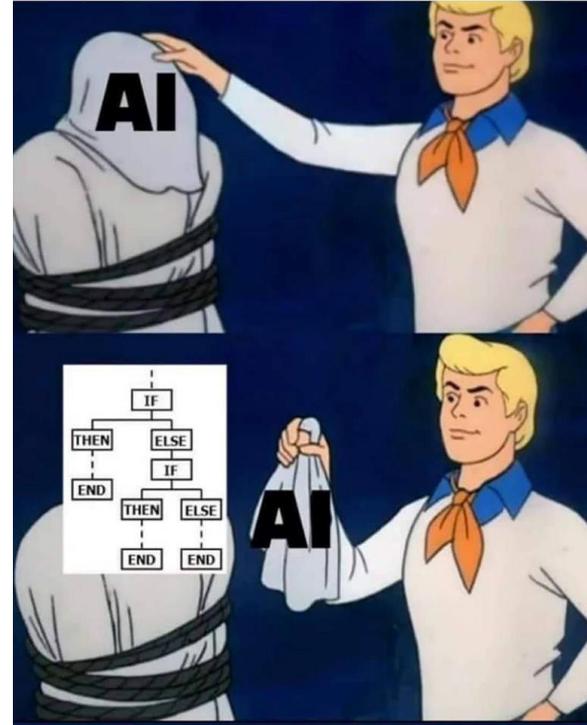


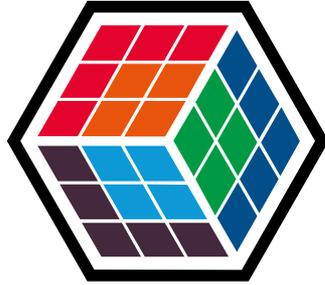
THE
DEVELOPER'S
CONFERENCE

Obrigado!



THE DEVELOPER'S
CONFERENCE





THE DEVELOPER'S CONFERENCE