



Speeding Up Machine Learning Development with MLflow

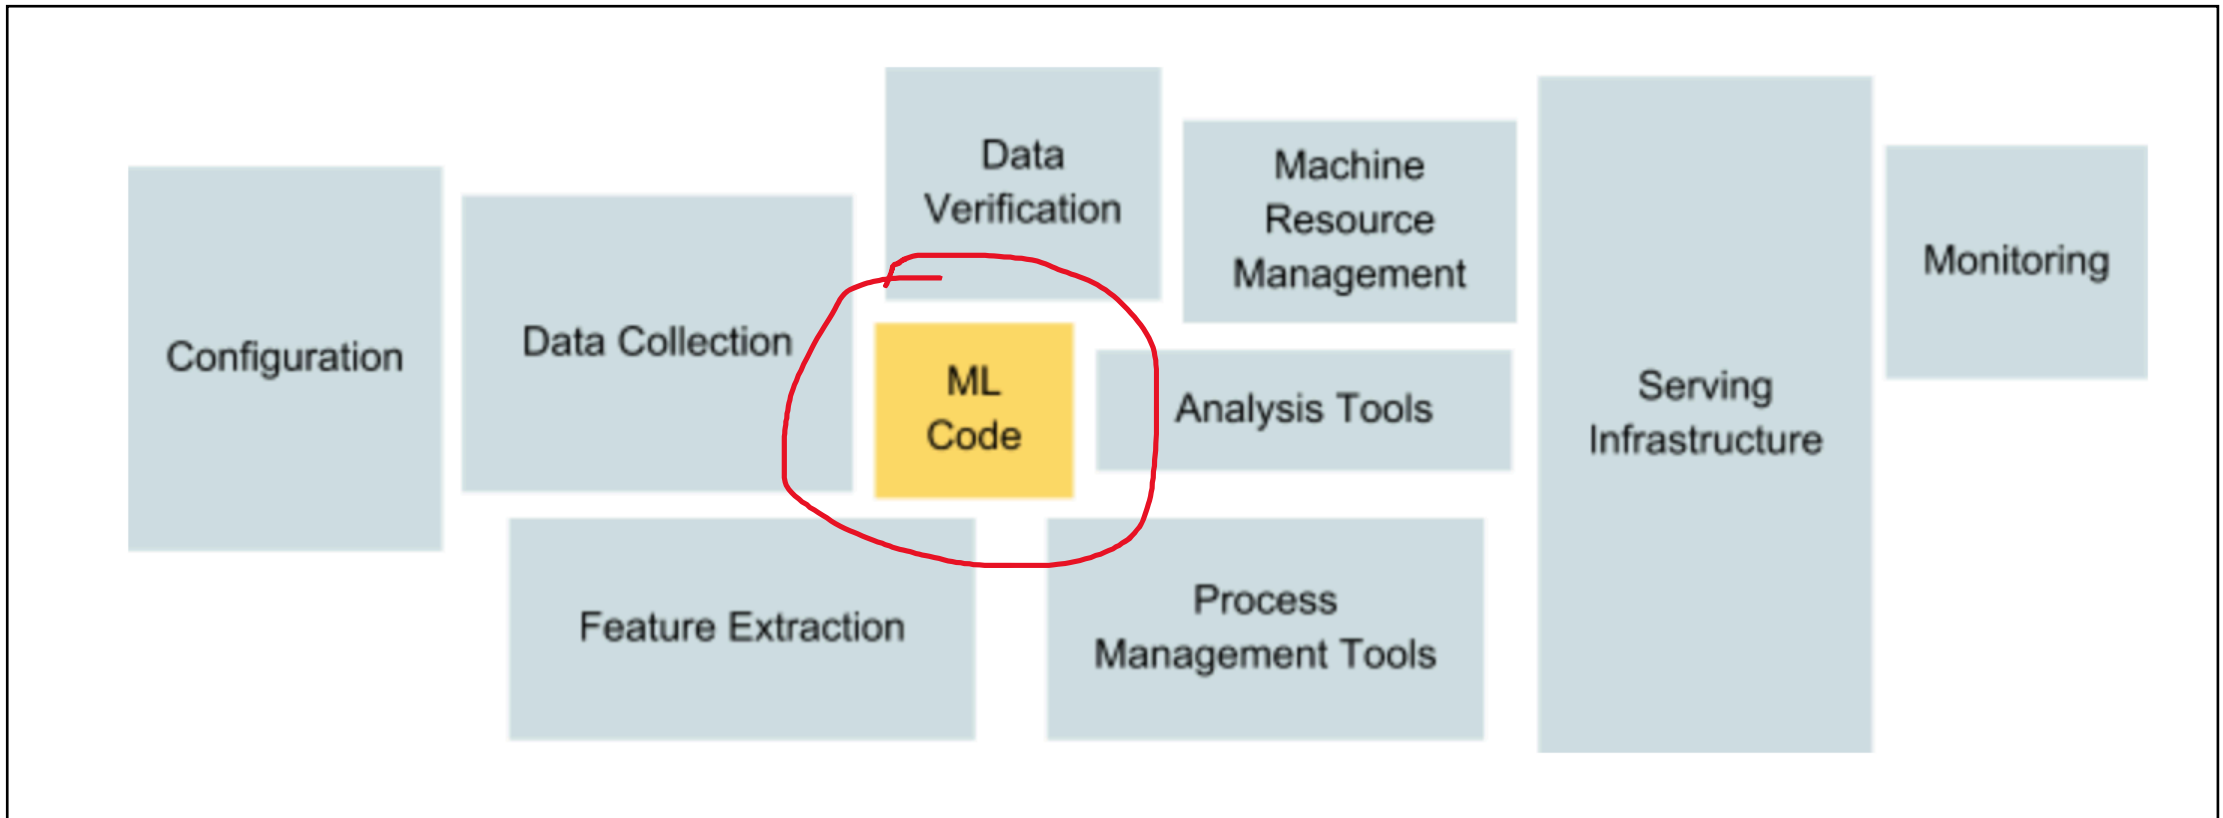
Hien Luu

Agenda

- Unique Challenges in ML Development Process
- Machine Learning Platform Tour
- Introduction to MLflow
 - Demo

ML Development Process Overview

[Hidden Technical Debt in Machine Learning Systems](#) (paper from Google)



The required surrounding infrastructure is vast & complex

ML Development Process Overview

Software 2.0

Data
Driven



Model
Driven

New dawn of a new age of software

ML Development Process Overview

Software vs Model Development

Software

Model

Goal

Meet function requirements

Optimize for business metric

Quality

Depend on code

Depend on data, code, algorithm, parameters

Tool

Standardized stack

Many tools and libraries

Outcome

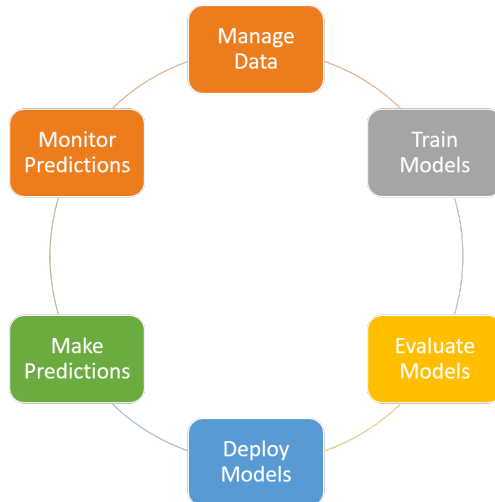
Fairly deterministic

Changing with data


ML Development Process Overview

Machine Learning Development Dimensions

Experimentation



Environments

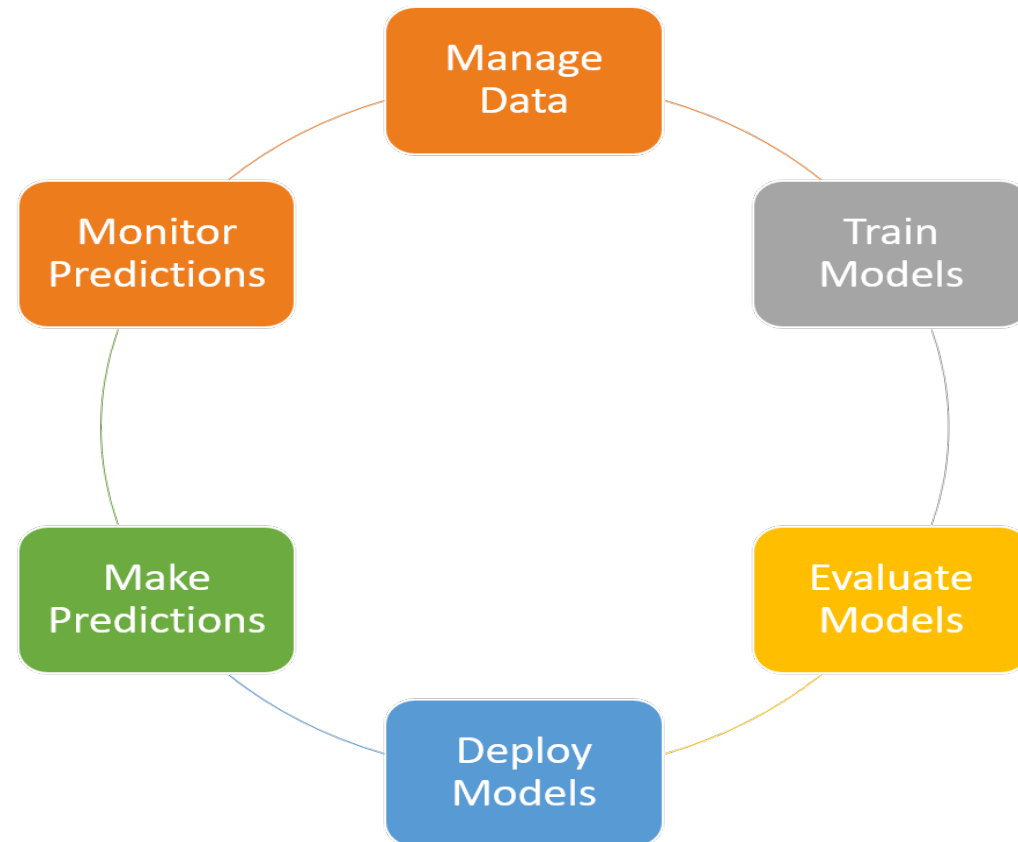
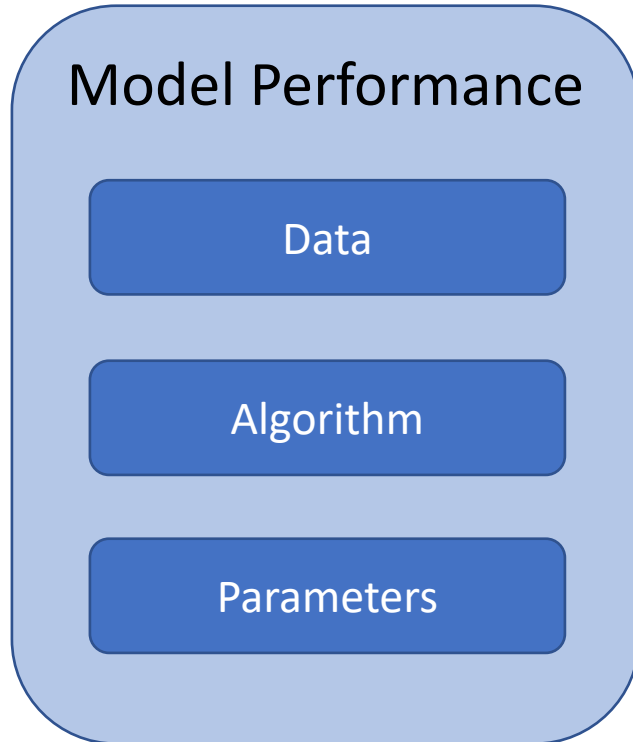


Big Data



ML Development Process Overview

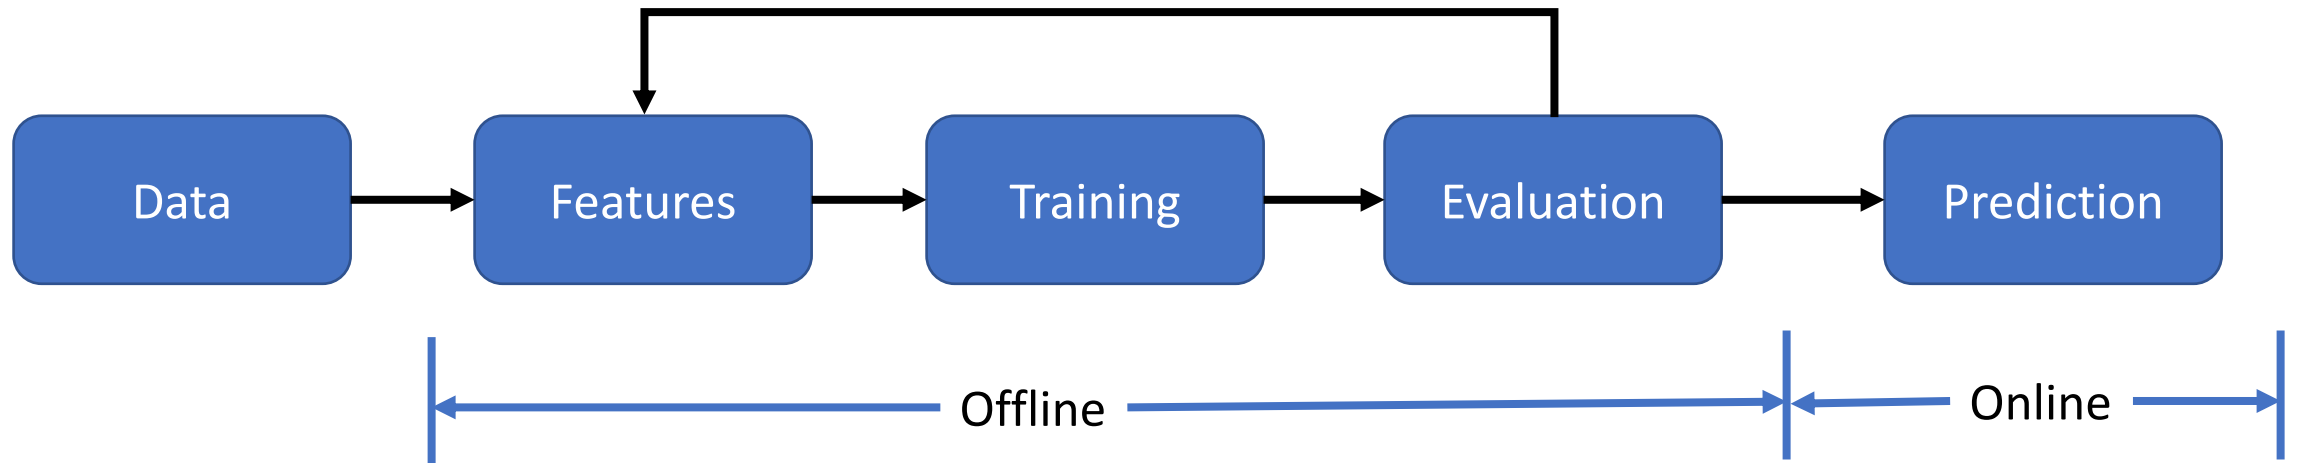
Experimentation Dimension



Machine Learning development is a scientific endeavor

ML Development Process Overview

Multiple Environments Dimension



ML Development Process Overview

Data Volume Dimension




Machine Learning Pipeline & ML Training Infrastructure

Machine Learning Platform Tour

Other Challenges Moving To Model Driven



People



Process

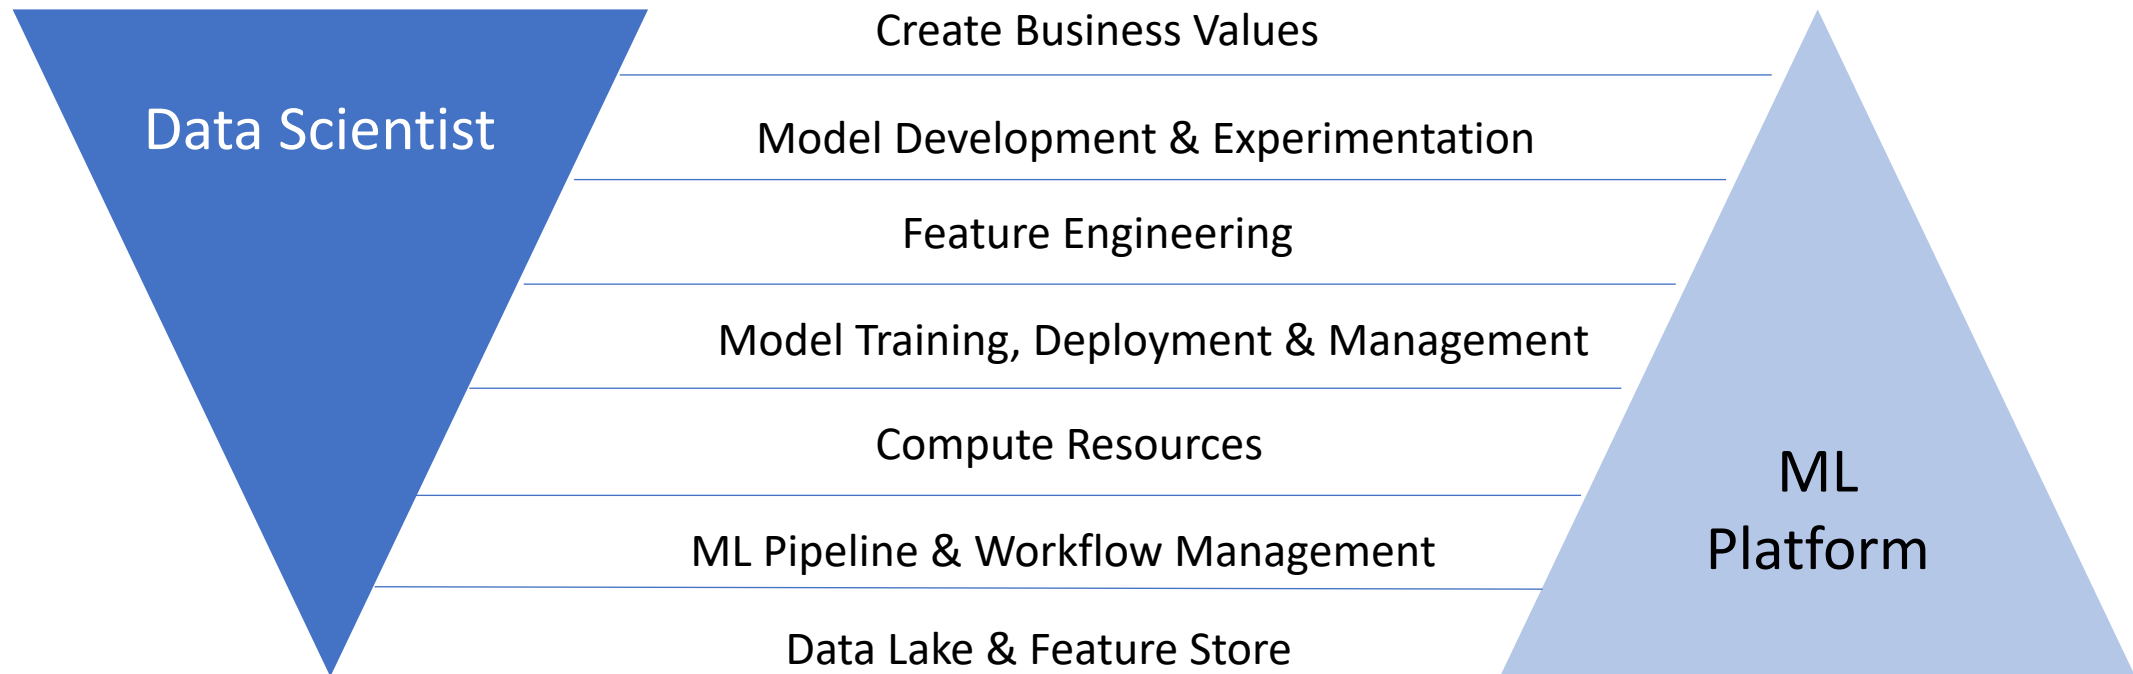


Technology

Scientific exploration and engineering rigorous and automation

ML Development Process Overview

Separation of Concerns

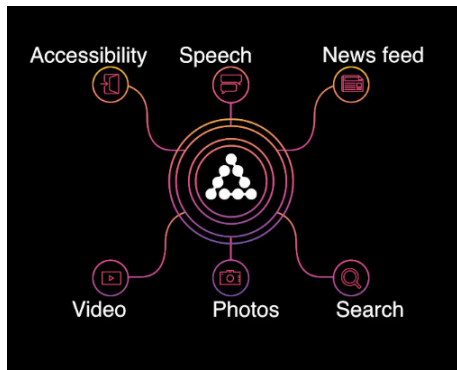


ML Development Process Overview

ML development vs traditional software
development

Machine Learning Platform Tour

In-house Machine Learning Platforms



FBLearn

AI Backbone

UBER

Michelangelo

ML as-a-service



Google TFX

*Deploying production
ML pipelines*



ML as-a-service

Automation, productivity and fast iteration

Machine Learning Platform Tour

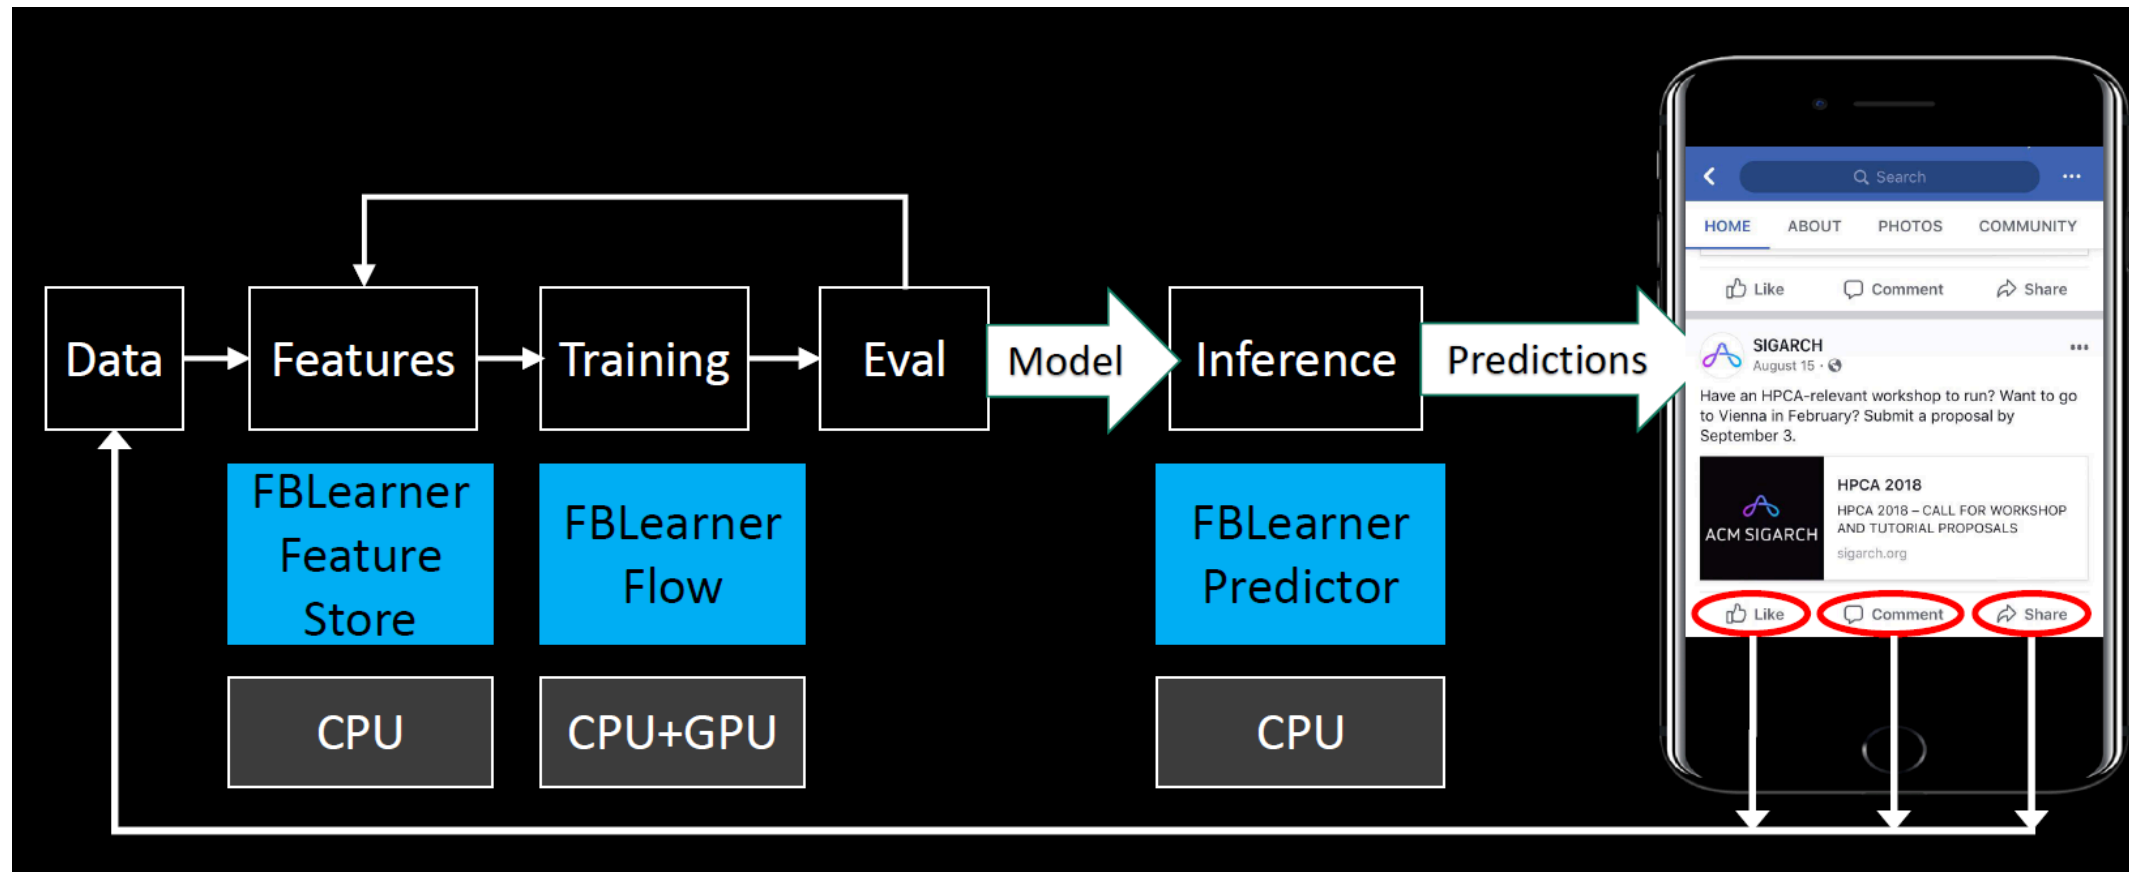
Anatomy of Machine Learning Platform



Minimize incidental complexity in Machine Learning to increase efficiency

Machine Learning Platform Tour

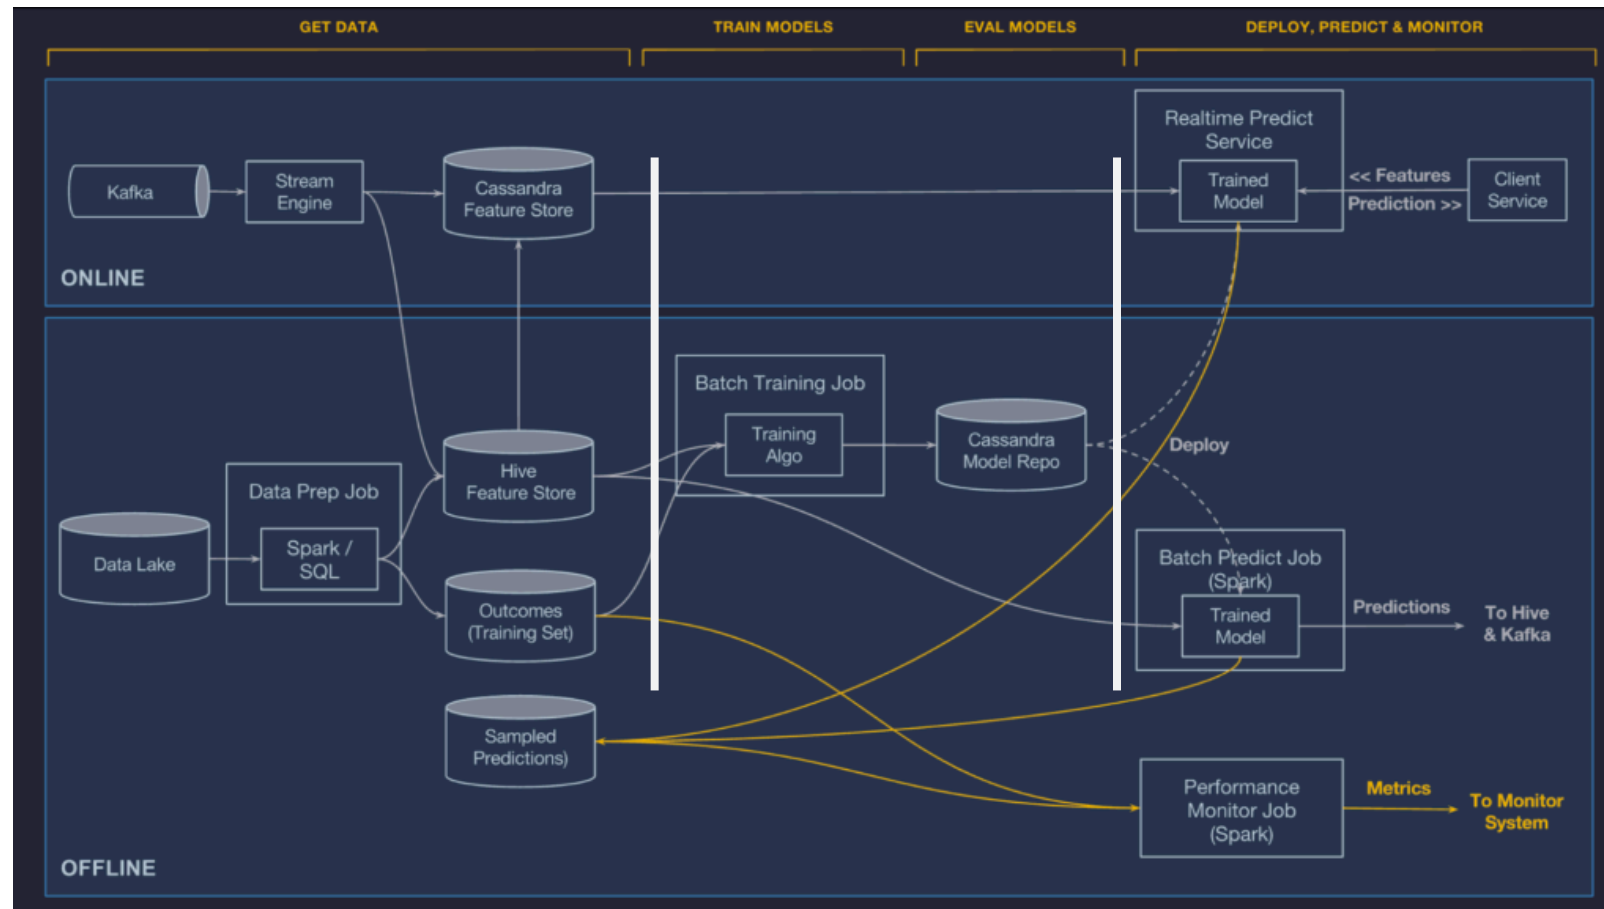
Facebook - FBLeaer



Productivity, reusability, scalability, and ease-of-use

Machine Learning Platform Tour

Uber - Michelangelo



Democratize & scale AI to make it as easy as requesting a ride

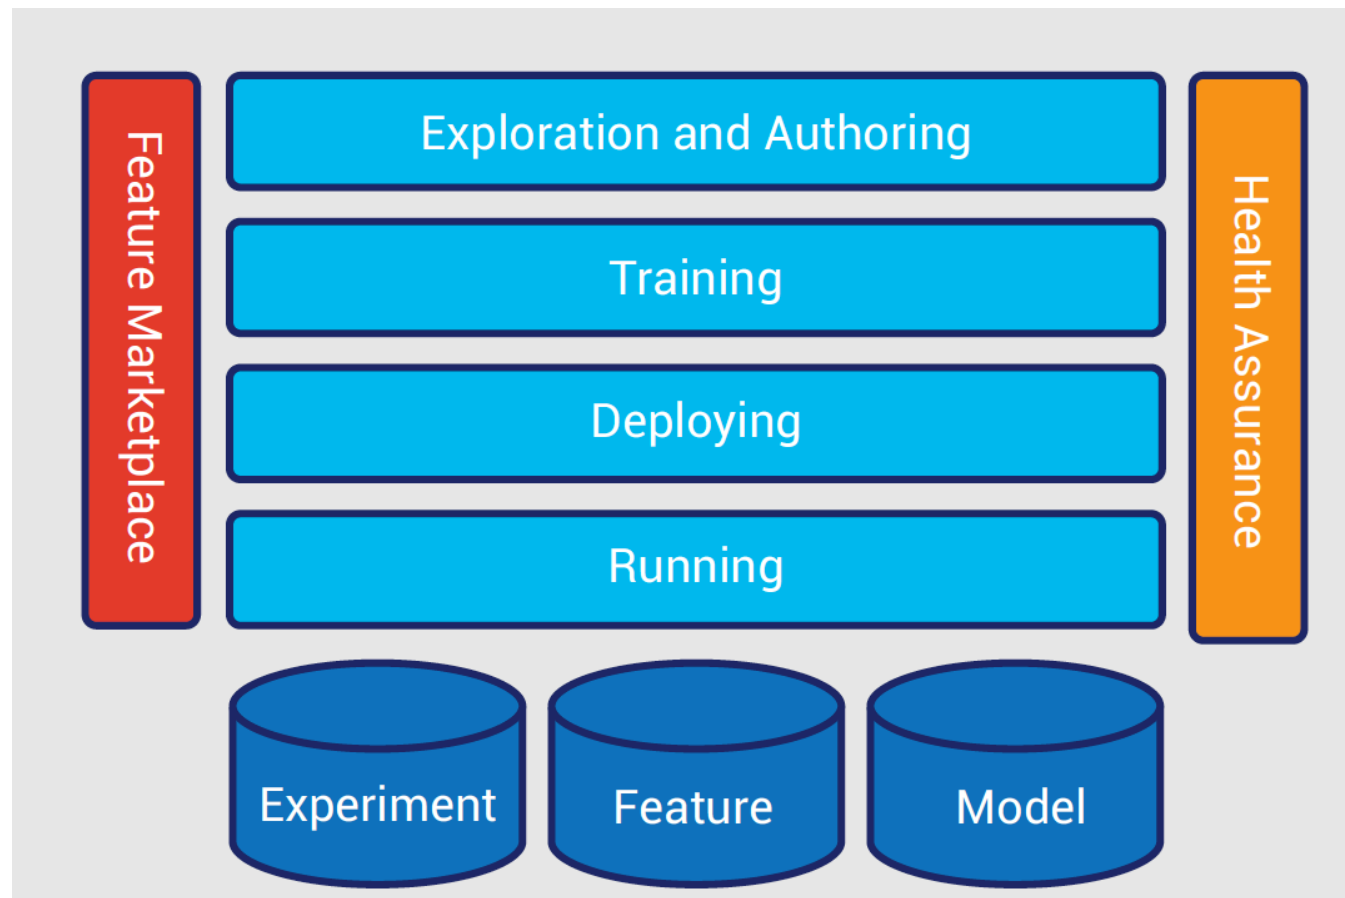
ML as software engineering

Iterative, tested, and methodical

Machine Learning Platform Tour

LinkedIn - Pro-ML

To double effectiveness of ML engineers



Use Cases

- Feeds
- Search
- Recommendation
- Advertisement
- Fraud

Machine Learning Platform Tour

Cloud Based Machine Learning Platforms

Machine Learning as a Service - MLaaS



AWS SageMaker



Azure ML



Google Cloud
Machine Learning Engine

Major pain points associated with a machine learning project dramatically change as the scale of the project increases.

Machine Learning Platform Tour

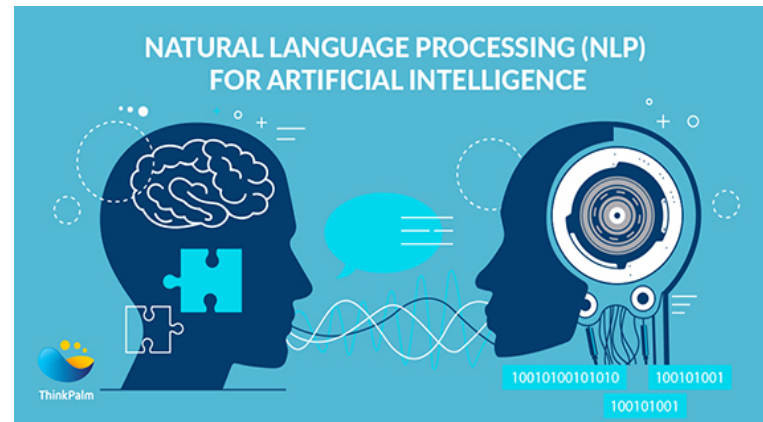
Pre-trained Machine Learning Models

Vision



Text within image, facial expressions

Speech



Chat bots, disease predictions,
fake news

Language



Translation, language detection

Machine Learning Platform Tour

Cloud MLaaS

AI Services

(Computer vision, object recognition, NLP)

ML Services

(ML IDE, experimentation, model training, management & monitoring)

ML Frameworks & Compute Infrastructure

(Tensorflow, Pytorch, Caffe, GPUs, Kubernetes, prediction infrastructure)

MLflow – Open ML Platform



Principles

- Open
- Ease of use
- Extensible
- Scalable

An open source platform for the
machine learning lifecycle

Manage the ML lifecycle, including experimentation, reproducibility and deployment

MLflow – Open ML Platform

mlflow

The MLflow logo is centered at the top. Below it, four blue boxes represent the core components of the platform. Arrows from each of these boxes point upwards towards the 'flow' part of the MLflow logo.

Tracking

Record and query experiments
(code, configs,
data,result)

Project

Packaging format for
reproducible runs on
various platforms

Model

Model format that
standardize
deployment

Model Registry

Model lifecycle
management

MLflow – Tracking

The MLflow logo is displayed in a stylized font. 'ml' is in white, 'f' is in blue, and 'low' is in white with a blue outline. The logo is centered within a dark blue rounded square.

mlflow

Tracking

Record and query
experiments: code,
configs, results, ...etc

Track and analyze experiments

MLflow – Tracking

Tracking Experiments

1	Dataset	Split (train/dev/test)	0.7/0.2/0.1	0.7/0.2/0.1	0.7/0.2/0.1
2		Class ratio (train/dev/test)	0.42/0.42/0.42	0.42/0.42/0.42	0.42/0.42/0.42
3		train/dev/test size	4871/1392/696	4871/1392/696	5315/1500/696
4	Training hyperparameters	Learning rate	1.00E-05	1.00E-05	1.00E-05
5		epoch	3	2	3
6		batch size	32	32	32
7	Results	accuracy	0.88304595	0.8650862069	0.8650862069
8		f1	0.82495437	0.8108753316	0.8108753316
9		precision	0.878865	0.7848381601	0.7848381601
10		recall	0.7780239	0.8389705882	0.807602069
11		tp	1398	1402	1460
12		tn	1692	1663	1707
13		fp	1113	1142	1161
14		fn	1189	1185	1190
15		loss	0.59637538	0.594134	0.594134
16		Test results	accuracy	0.90747	0.90747
17		f1	0.85636	0.85636	0.83108
18		precision	0.90934	0.90934	0.86689
19		recall	0.8099	0.8099	0.79846
20					
21					
22					
23					
24					
25					
26					
27					
28					
29					
30					

“ML experimentation is like the wild west. Ad-hoc tools and processes because of a lack of standardized tooling. Forget reproducibility, it is difficult to track experiments and results.”

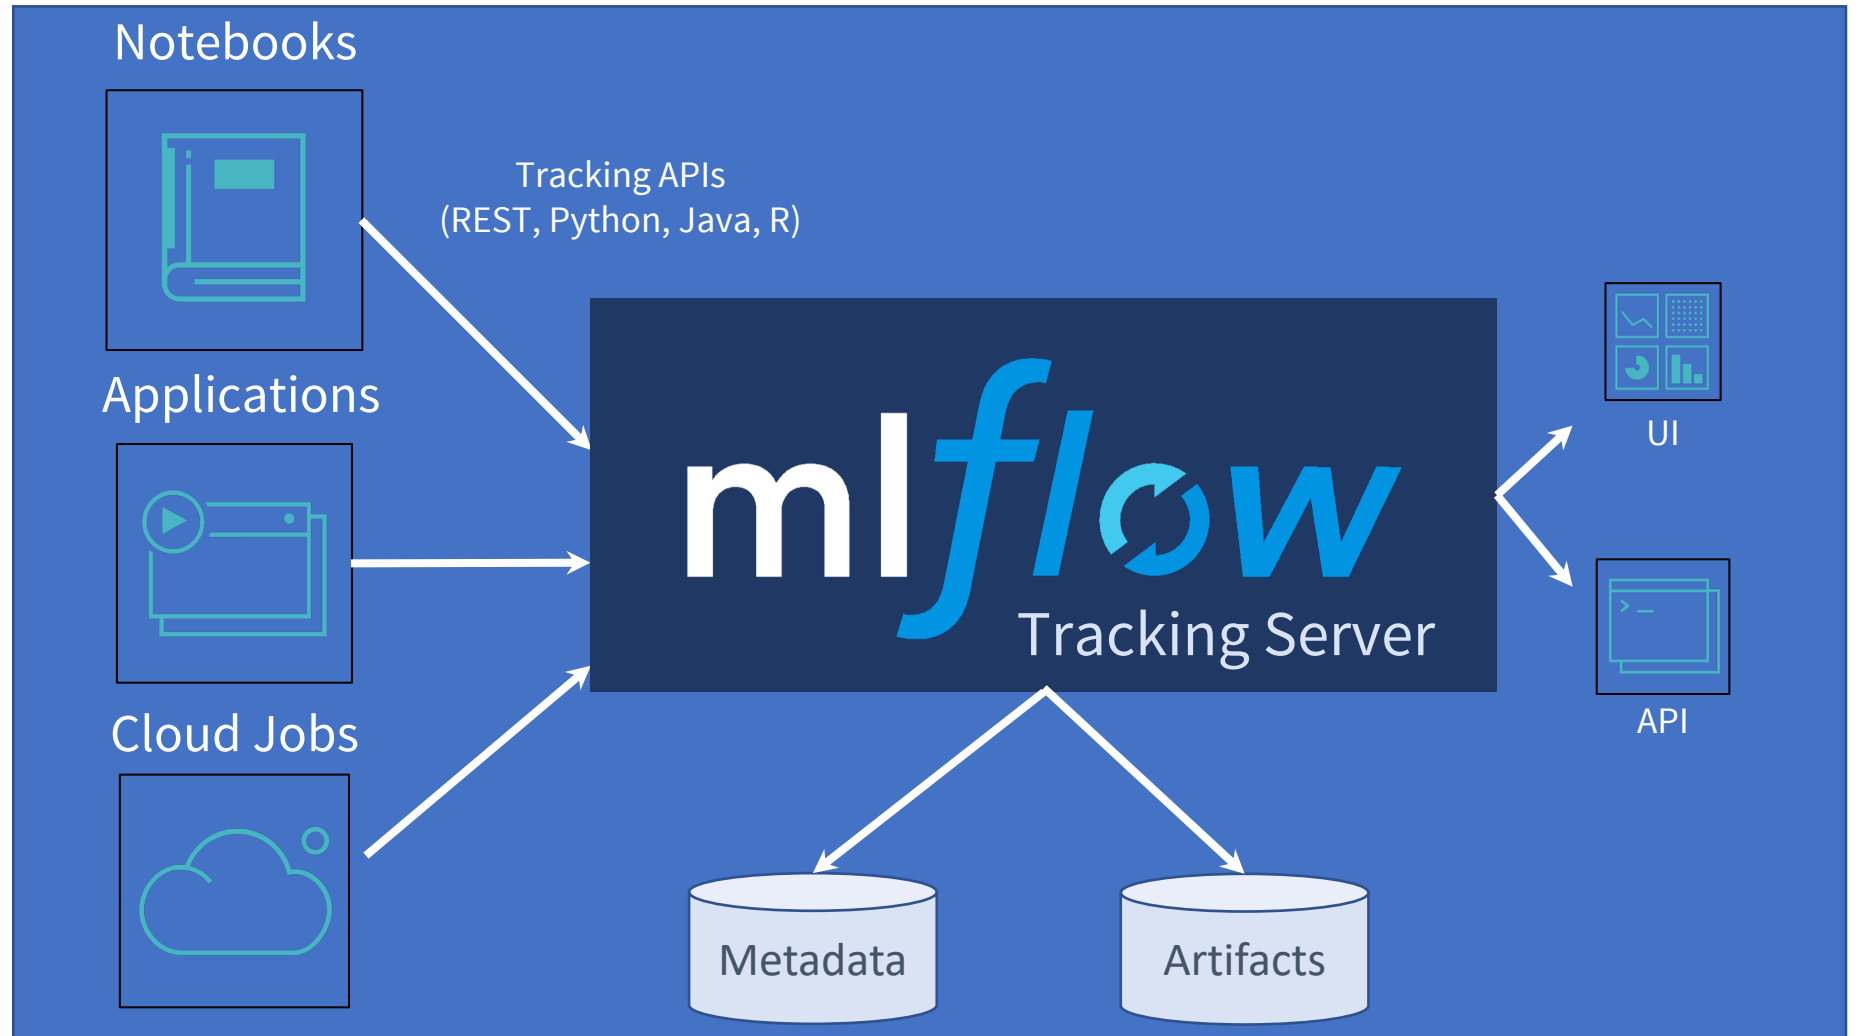
MLflow – Tracking

Tracking

Record and query experiments: code, data, config, results

Key Concepts

- K-V parameters
- Metrics
- Artifacts
- Source code
- Version (git)
- Tags & notes



MLflow – Tracking

Python, R, Java, REST

Tracking

Record and query experiments: code, data, config, results

```
import mlflow

# log model's tuning parameters

with mlflow.start_run():
    mlflow.log_param("layers", layers)
    mlflow.log_param("alpha", alpha)

# log model's metrics
mlflow.log_metric("mse", model.mse())
mlflow.log_artifact("plot", model.plot(test_df))
```

Recently added: `mlflow.keras.autolog()`

MLflow – Open ML Platform

- UI Demo
 - Show MLFlow parameter and metric logging
 - Visualize the experiments
 - `mlflow ui`
 - <http://127.0.0.1:7000/#/>

MLflow – Projects

The logo for MLflow Projects is centered within a dark blue rounded square. It features the word "mlflow" in a lowercase, sans-serif font, with "ml" in white and "flow" in blue. Below it, the word "Projects" is written in a white, sans-serif font.

mlflow

Projects

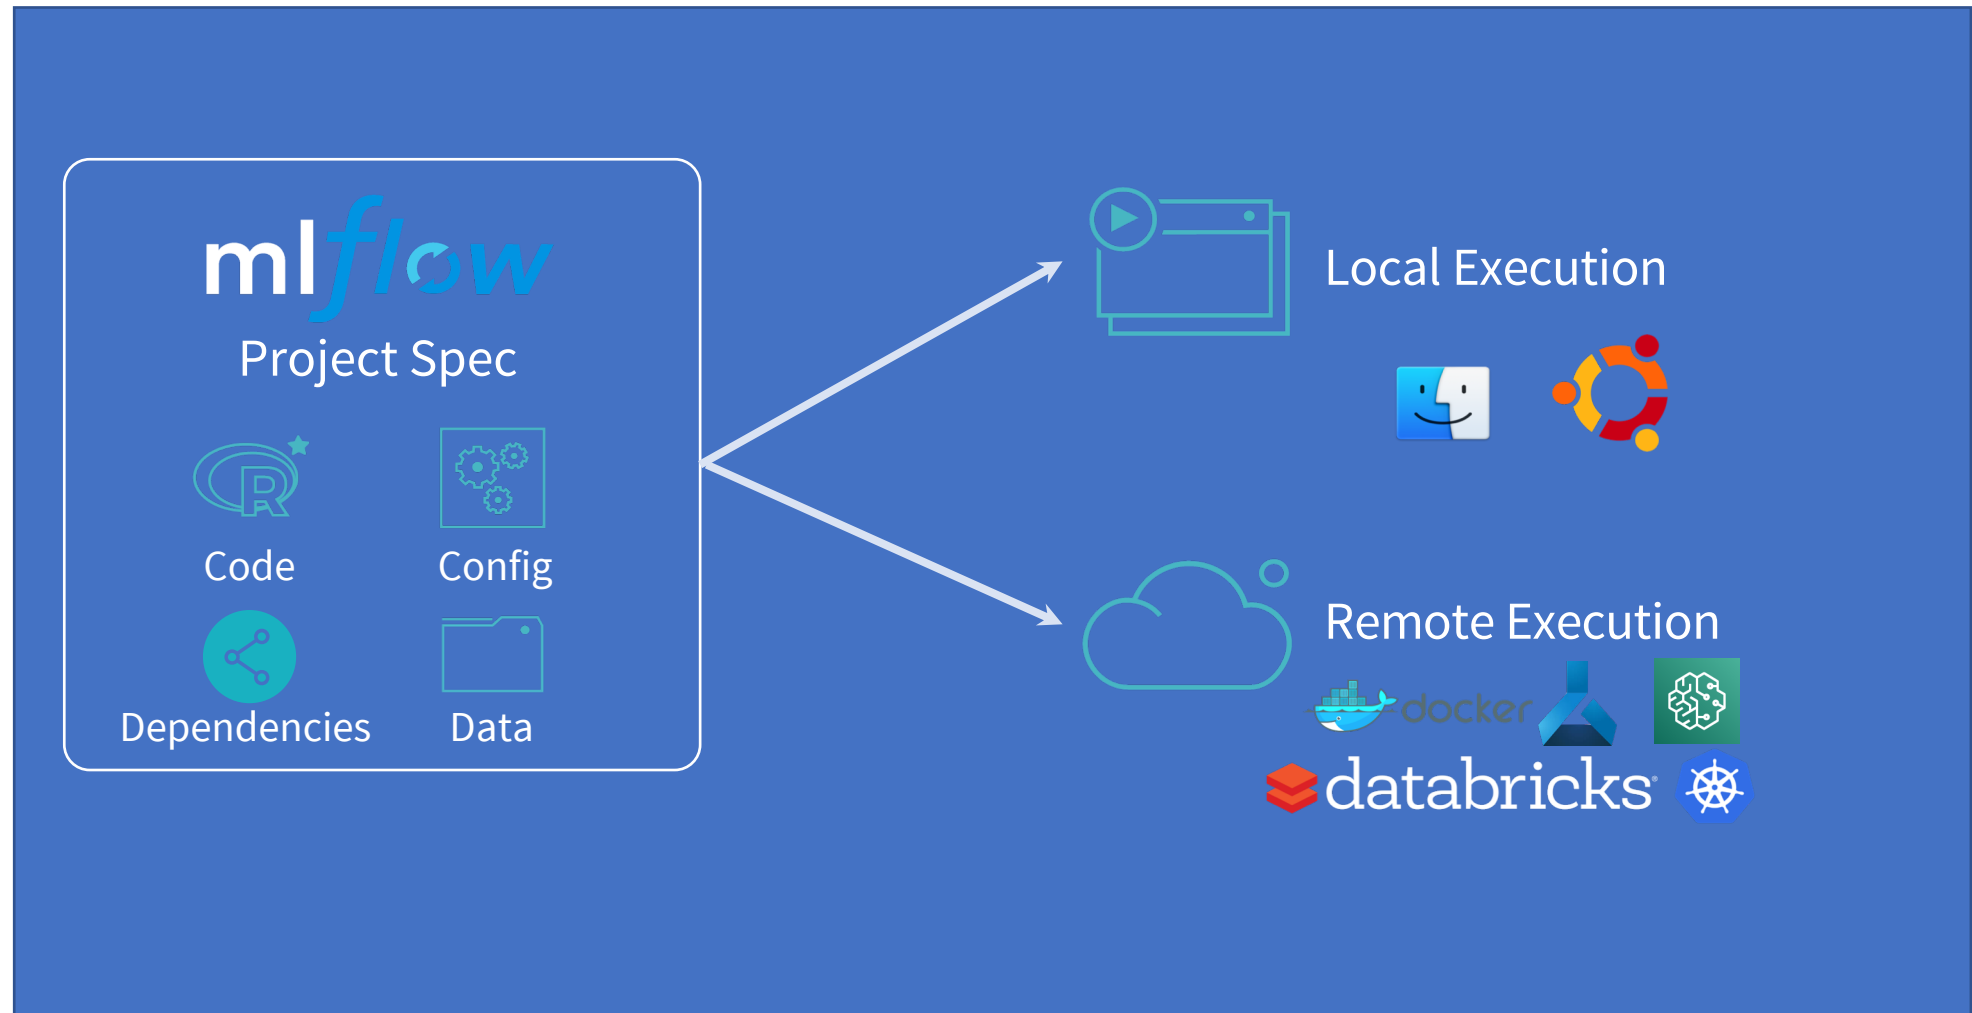
Packaging format
for reproducible runs on any
platform

Reproducibility via self-contained ML project specification

MLflow – Open ML Platform

Projects

Packaging format for reproducible runs on any platform



Reproducibility, Sharing, Productionalization

MLflow – Projects

Projects

Packaging format for
reproducible runs
on any platform

```
my_project/  
├── MLproject  
  
├── conda.yaml  
├── main.py  
├── model.py  
└── ...
```

```
conda_env: conda.yaml
```

```
entry_points:
```

```
  main:
```

```
    parameters:
```

```
      training_data: path
```

```
      lambda: {type: float, default: 0.1}
```

```
    command: python main.py {training_data}  
             {lambda}
```

```
$ mlflow run <directory> or git://<my_project>  
mlflow.run("<directory> or git://<my_project>")
```

MLflow – Projects

- Demo
 - `mlflow run <local directory>`
 - `mlflow run <github>`

MLflow – Models

The MLflow logo is displayed in a stylized font. 'ml' is in white, 'f' is in blue, and 'low' is in white with a blue outline. The logo is centered within a dark blue rounded square.

mlflow

Models

General model format
that supports diverse
deployment tools

Simplify model deployment

MLflow – Models

Models


General format for sending models to diverse deploy tools

ML Frameworks




mlflow
Model Format

Flavor 1



Flavor 2



Standard for ML models

Serving Tools



Inference Code



Batch & Stream Scoring



MLflow – Models

Models

General format for sending models to diverse deploy tools

```
mlflow.tensorflow.log_model(...)
```

```
my_model/  
├── MLmodel  
└── estimator/  
    ├── saved_model.pb  
    ├── variables/  
    └── ...
```

```
run_id: <uuid>  
time_created: 2019-06-20T08:11  
flavors:  
  tensorflow:  
    saved_model_dir: estimator  
    signature_def_key: predict  
python_function:  
  loader_module:  
    mlflow.tensorflow
```

MLflow – Models

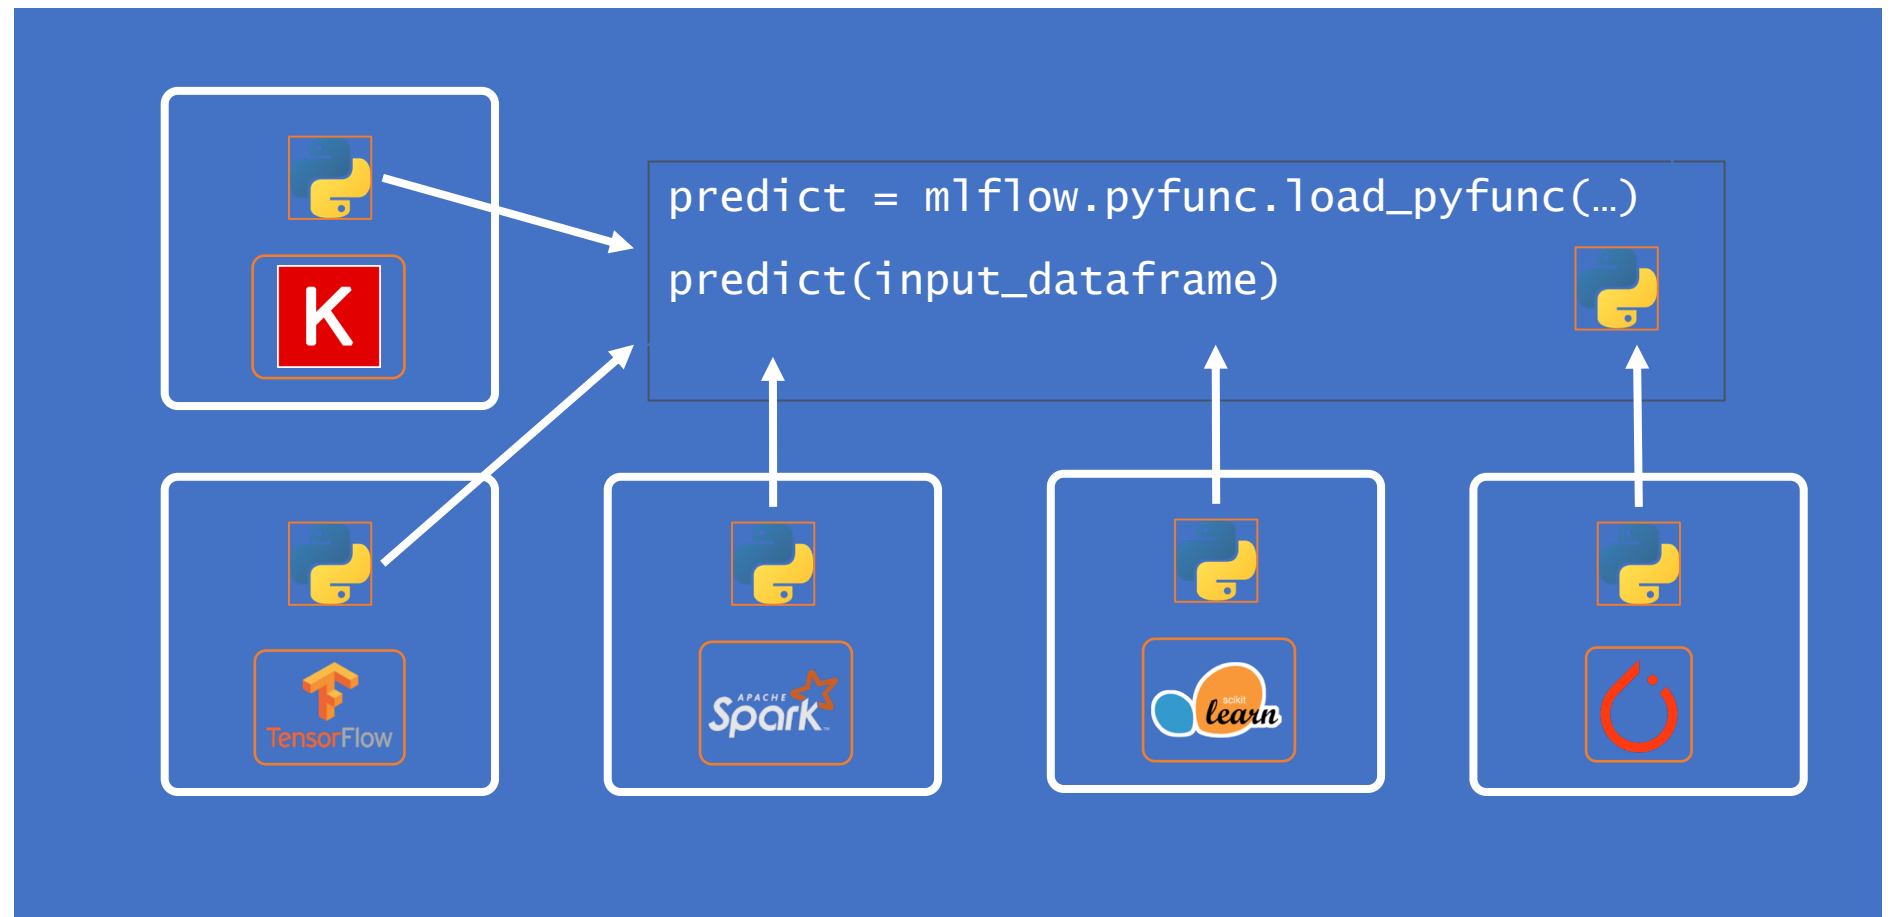
Model Flavors

Models

General format for sending models to diverse deploy tools

Built-In Flavors

- Python Function
- R,H₂O
- Keras
- MLeap
- PyTorch
- Skiki-learn
- Spark Mlib
- TensorFlow
- ONNX



MLflow – Model Registry

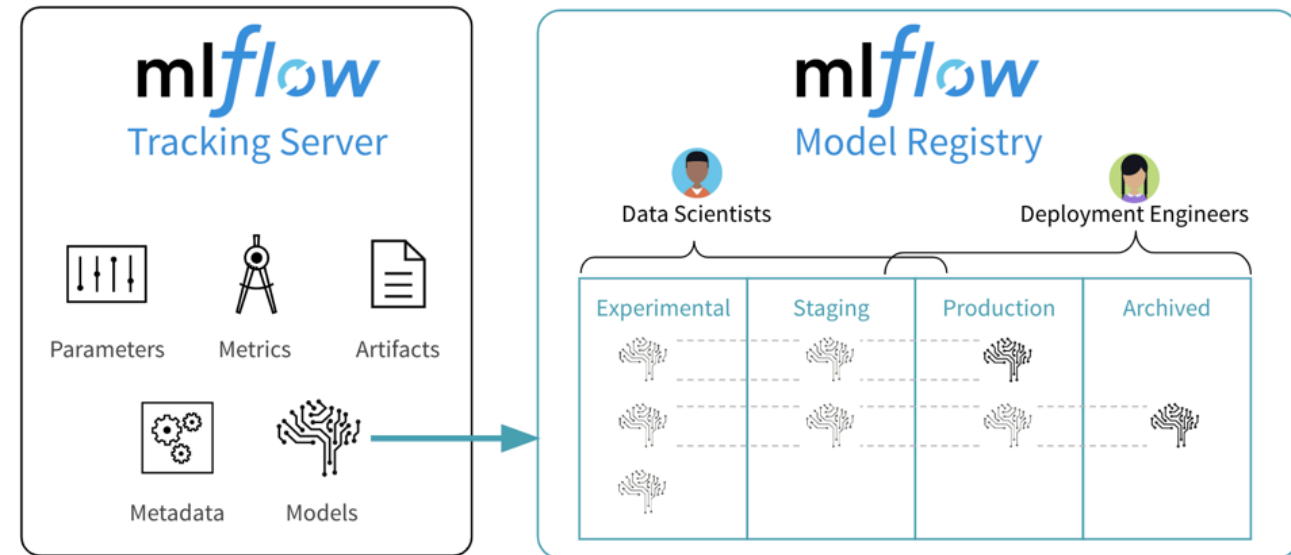


Collaboratively manage the full lifecycle of a model

MLflow – Model Registry

Managing Models Collaboratively

- Model administration and review
 - Sharing, versioning, approval workflow
- Integrate with MLflow tracking
- Centralized activity logs and comments
- Integration with CI/CD



MLflow

- Demo
 - Deploy model to a REST endpoint
 - `mlflow models serve`
 - Simple application to call REST endpoint to perform prediction
 - Flask application

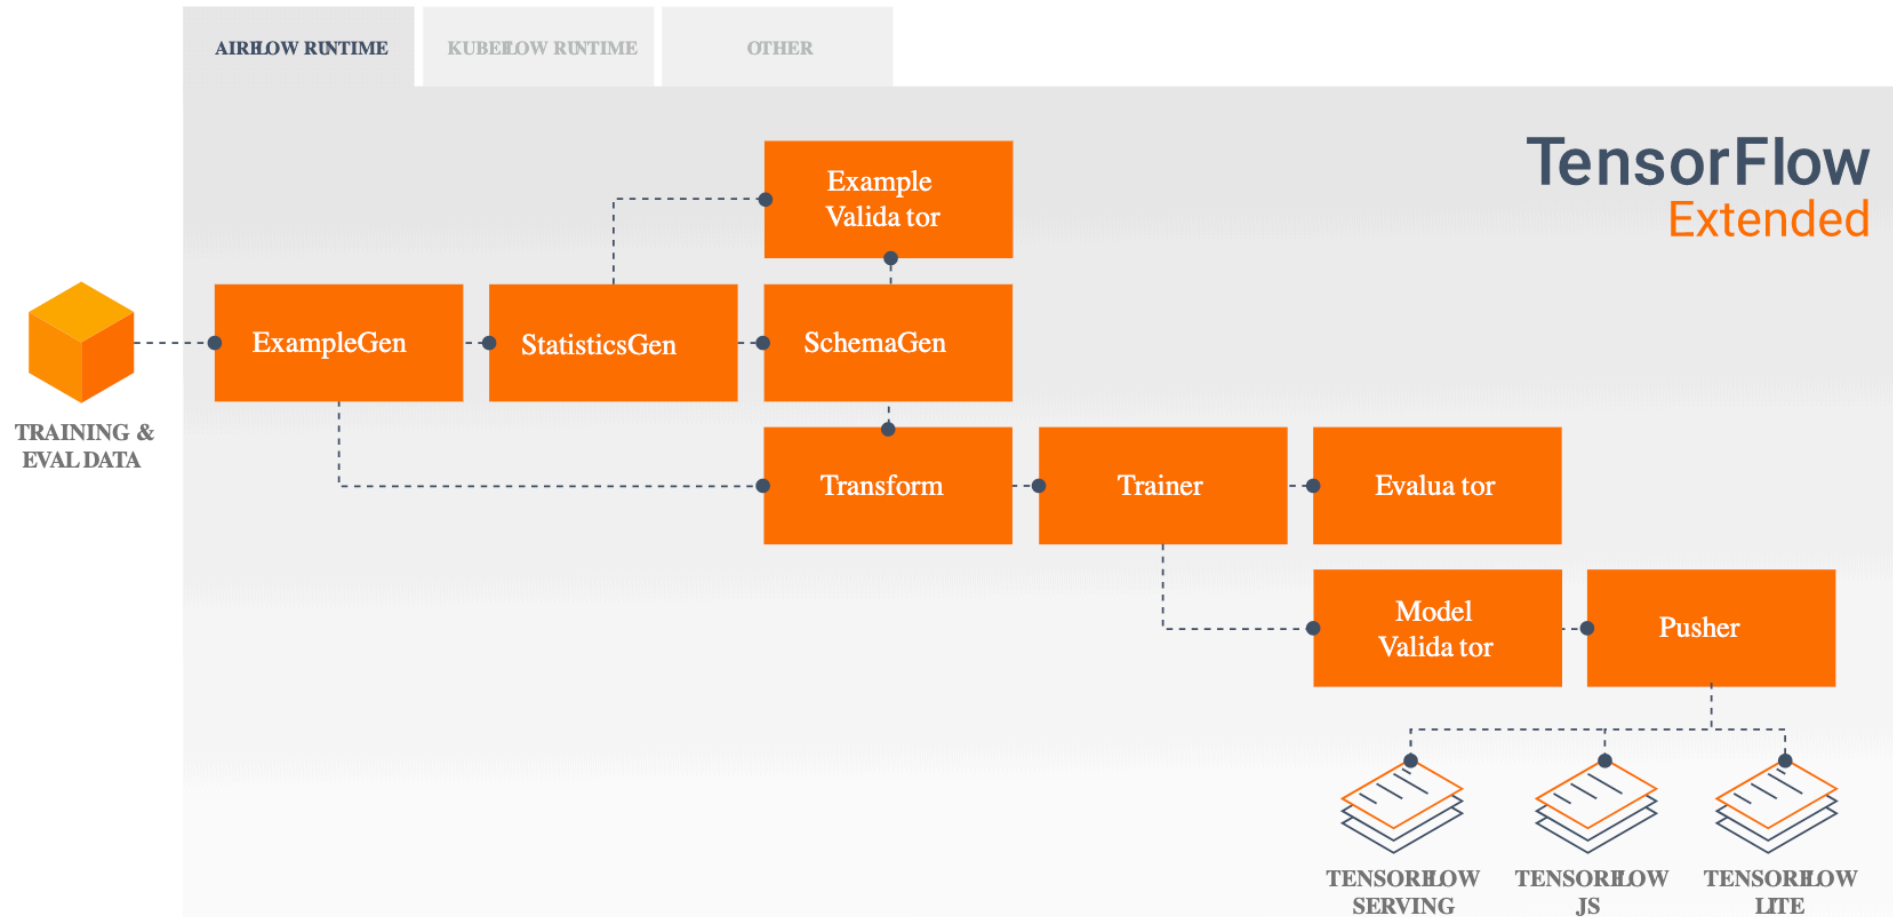
Summary

- Interesting & unique challenges with ML development process
- In-house & cloud-based ML infrastructures
- MLflow – open source ML platform
 - Tracking
 - Project
 - Model
 - Model registry

Q&A

Machine Learning Platform Tour

Google - TFX



ML Development Process Overview

People

Data
Engineer

Data
Scientist

Software
Engineer

Technologies

PYTORCH

TensorFlow

APACHE
Spark

dmlc
XGBoost

docker



databricks



Environments

Machine Learning Platform Tour

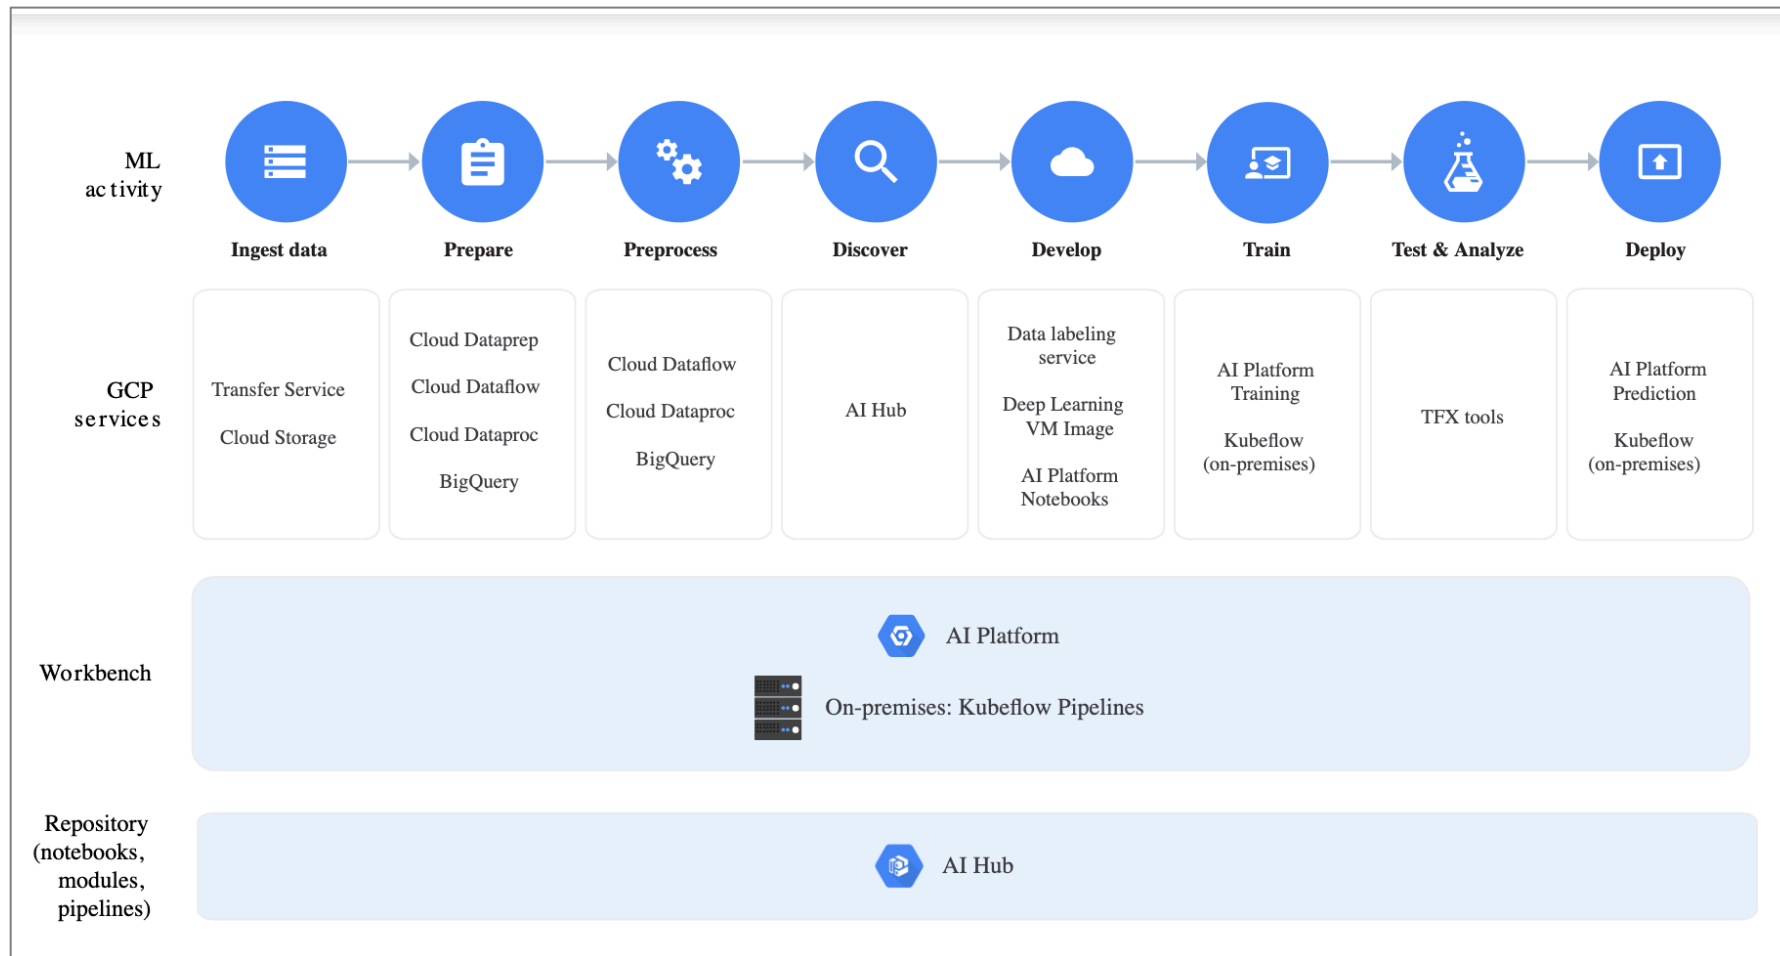
AWS AI Platform

The diagram illustrates the AWS AI Platform architecture, organized into three main sections:

- AI services:** A row of 12 service tiles. Vision includes Amazon Rekognition. Speech includes Amazon Polly and Amazon Transcribe (+Medical NEW). Text includes Amazon Comprehend (+Medical) and Amazon Translate. Search includes Amazon Kendra (NEW). Chatbots includes Amazon Lex. Personalization includes Amazon Personalize. Forecasting includes Amazon Forecast. Fraud includes Amazon Fraud Detector (NEW). Development includes Amazon CodeGuru (NEW). Contact centers includes Contact Lens For Amazon Connect (NEW).
- ML services:** A row of 10 service tiles. Amazon SageMaker is the primary service. Other tiles include Ground Truth, Augmented AI, ML Marketplace, Amazon SageMaker Studio IDE (NEW), and Neo. A white box highlights the SageMaker Studio IDE components: Built-in algorithms, Notebooks (NEW), Experiments (NEW), Model training & tuning, Debugger (NEW), Autopilot (NEW), Model hosting, and Model Monitor (NEW).
- ML frameworks & infrastructure:** A row of 10 infrastructure tiles. Frameworks include TensorFlow, mxnet, GLUON, and Keras. Other tiles include PYTORCH, Deep learning AMIs & containers, GPUs & CPUs, Amazon Elastic Inference, AWS Inferentia, and FPGA.

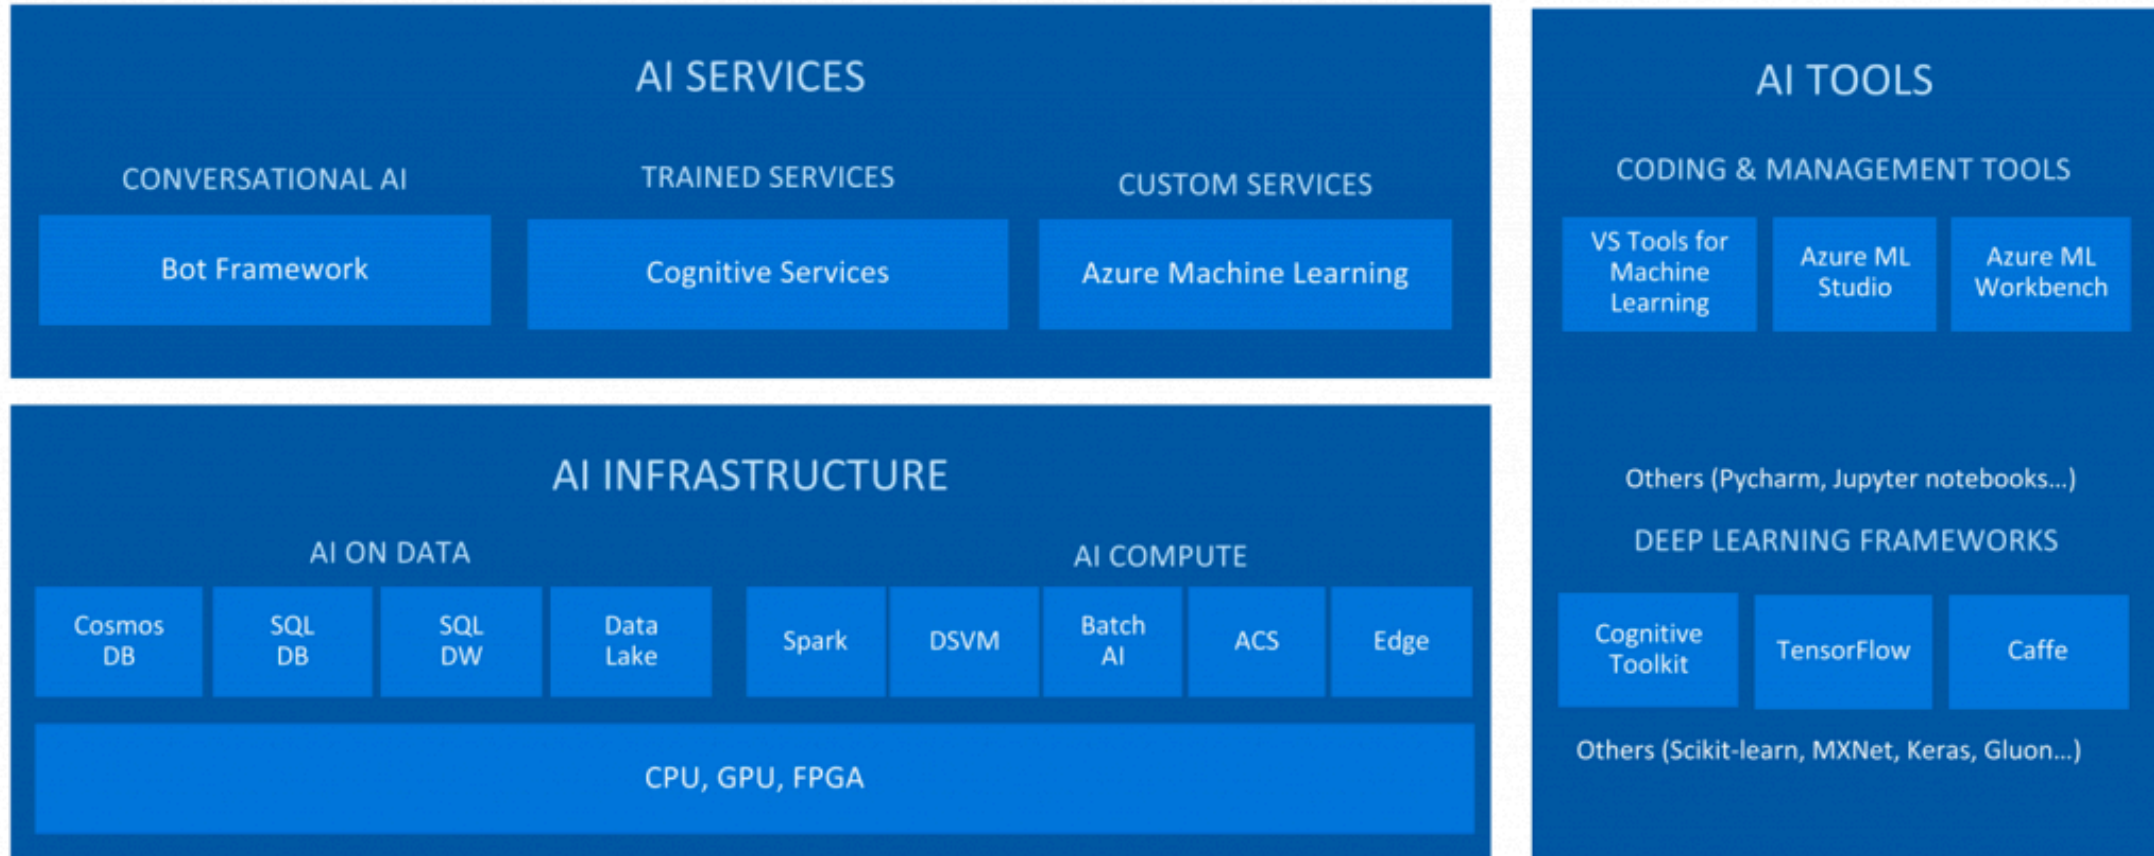
Machine Learning Platform Tour

Google AI Platform



Machine Learning Platform Tour

Azure AI Platform



MLflow – Models

Models

General format for sending models to diverse deploy tools

- Packaging format for ML Model
 - Directory structure
 - Different flavors
 - (keras, pytorch, sklearn, spark, tensorflow)
- Define dependencies
 - Easy reproducibility & deployment
- Model creation utilities
 - To save model in MLFlow format
- Deployment APIs
 - CLI – Python, R, Java

MLflow – Projects

Projects

Packaging format for
reproducible runs
on any platform

- Packaging format for reproducible for ML runs
 - Training code folder or Github repository
 - Project configuration file in YML format
- Define dependencies for reproducibility
 - Using Conda to specifying and managing dependencies
- Execution API for running projects
 - Local or remote execution
 - CLI – Python, R, Java
 - `mlflow run`