



Selecting Tools for Developing, Monitoring and Maintaining ML Models

Parshad Patel

Data Scientist

Yummly, Inc (part of Whirlpool)

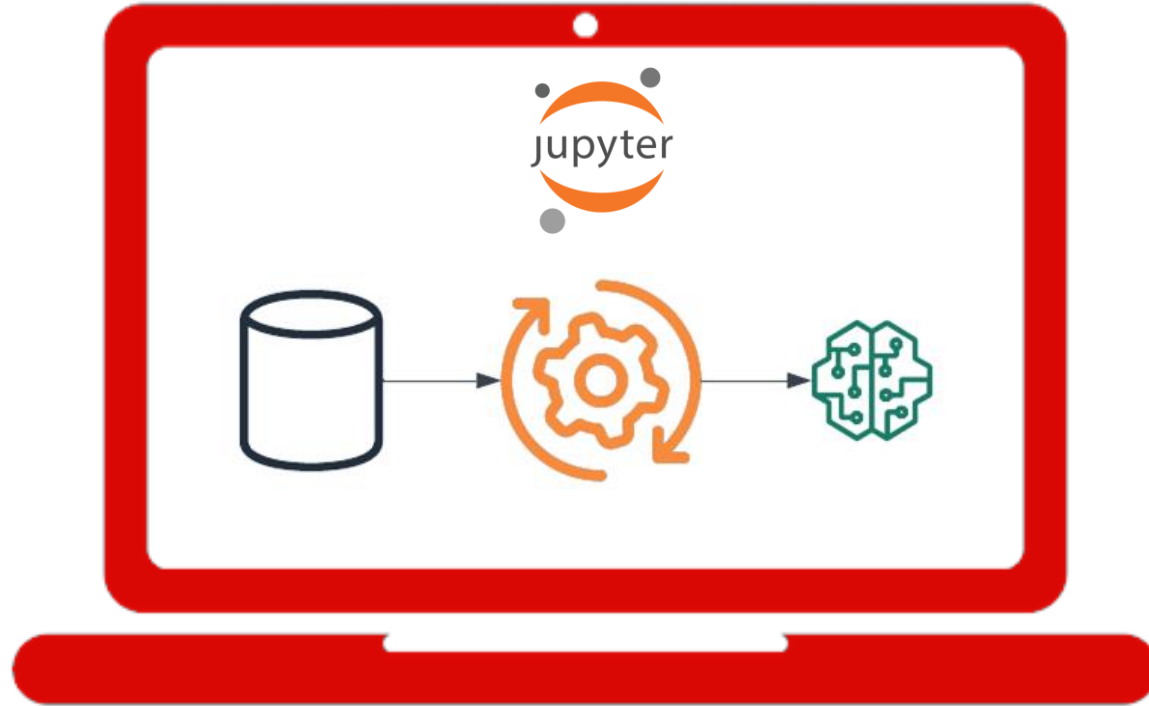


Agenda

- Navigating the landscape
- Building minimum viable pipelines
- Use-case at Yummly
- Takeaway



A data scientist's workflow



Navigating landscape - The paradigms

Feature Store

Stream Processing

Hyperparameter Optimization

Observability

CI/CD/CT

Visualization

Versioning



Batch Prediction

Inference

Labeling

Metadata Store

Tracking

AutoML

Model Registry

Navigating landscape - Tools and platforms



Start with first principles



Enter: Minimum viable pipelines




Minimum viable pipelines - Phase 1

Experiment tracking

Model registry

Business impact




Experiment tracking and model registry




Minimum viable pipelines - Phase 2

**Hyperparameter
tuning**



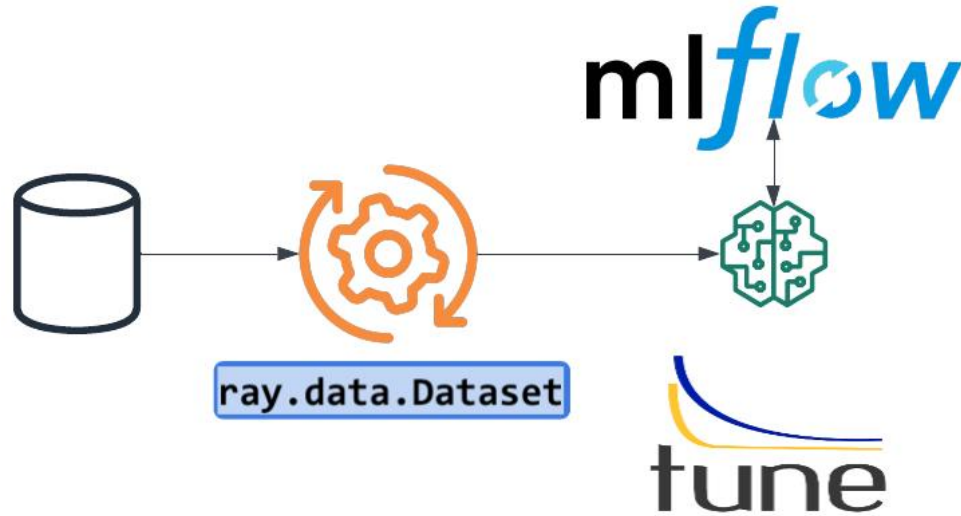
Data ingestion



Business impact




Hyperparameter tuning and data ingestion



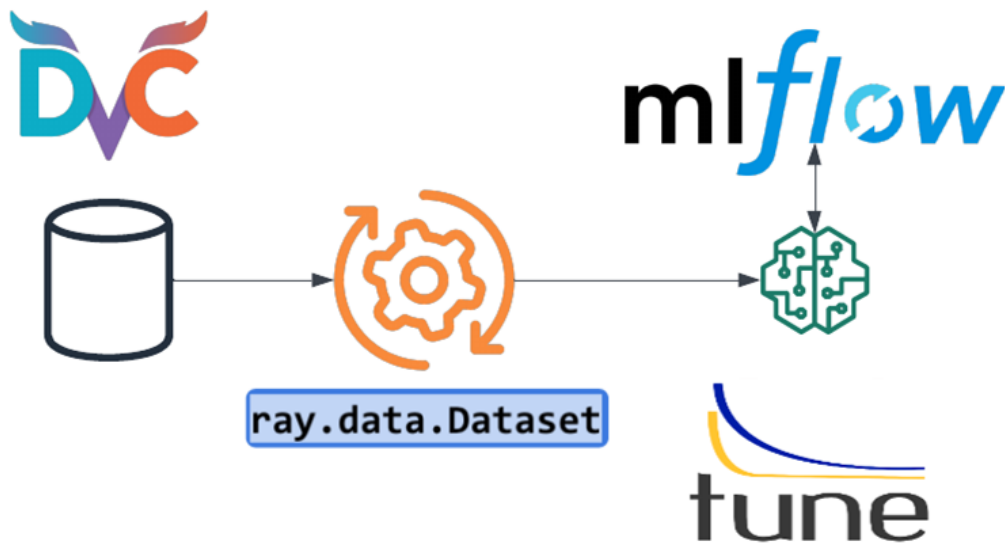
Minimum viable pipelines - Phase 3

**Dataset and artifact
tracking**

Business impact



Dataset and artifact tracking



Minimum viable pipelines - Together

mlflow




 RAY




 **DVC**


Minimum viable pipelines - Together



Minimum viable pipelines - Towards scaling





Use-case: Yummly



Use case - Yummly



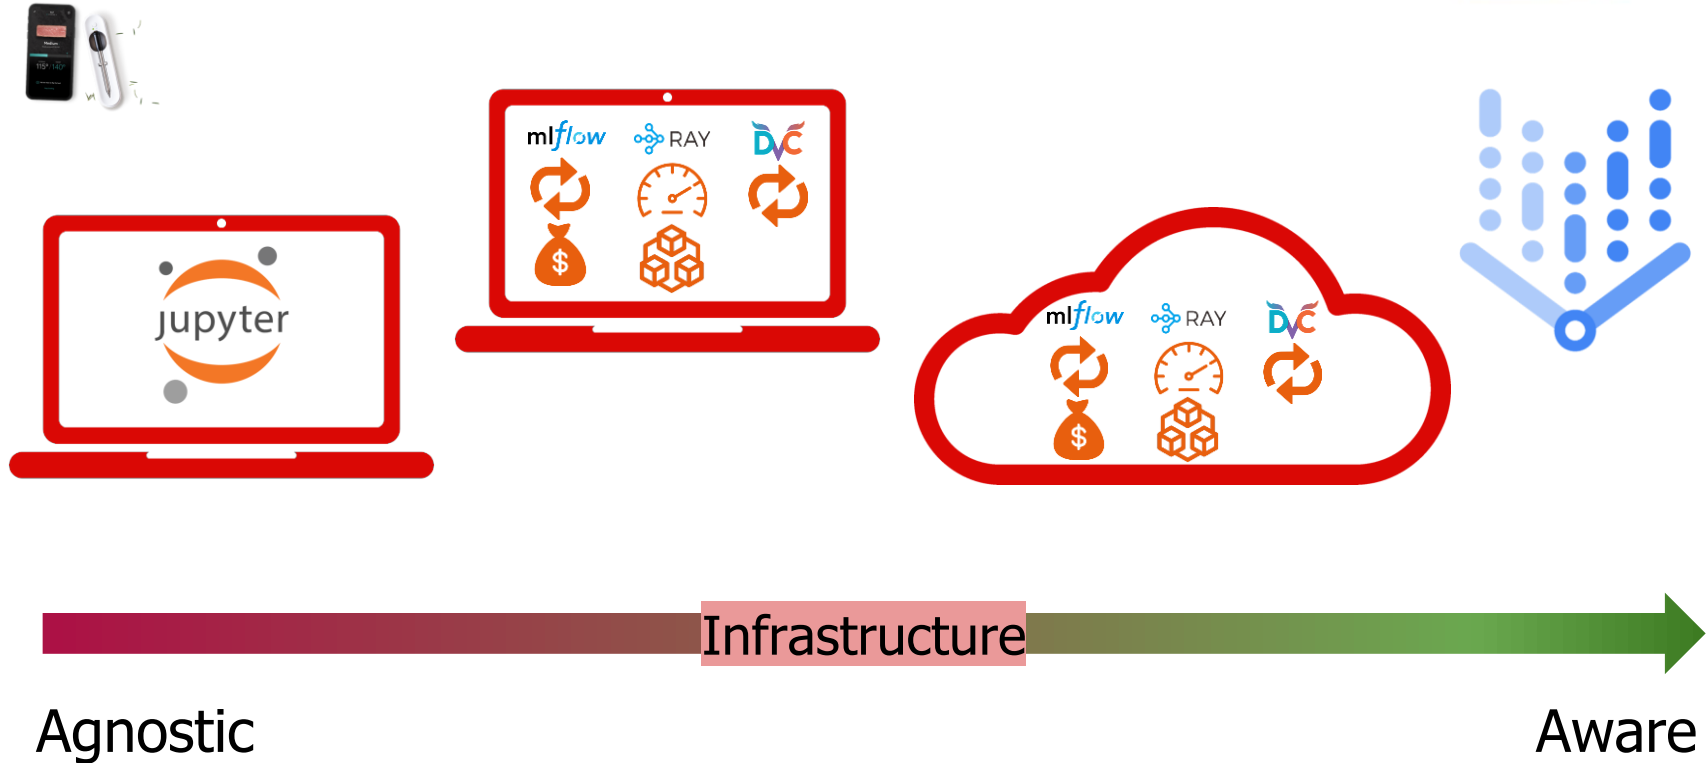
Use case - Journey



Minimum viable decision - Pause



Use case - Journey



Agnostic

Aware

Migration from build to buy

mlflow

 RAY

 DVC




Pipelines
Prediction
Metadata
Workbench
Feature Store
Matching Engine
Deep Learning Containers

Takeaway



In gist



In gist



Thank you!





[Ray Datasets](#)

[Ray Tune](#)

[Ray Train](#)

[MLFlow Tracking](#)

[MLFlow Registry](#)

[DVC Docs](#)

[Vertex AI Platform](#)