Introducing Oracle Machine Learning

A Collaborative Zeppelin notebook for Oracle's machine learning capabilities

Charlie Berger
Marcos Arancibia
Mark Hornick

Advanced Analytics and Machine Learning

ORACLE

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

What is a “notebook”

What is a “notebook”

- Web-based – accessible by browser
- Enables interactive data analytics
- Produces appealing data-driven and collaborative documents
- Integrates formatted notes with your code

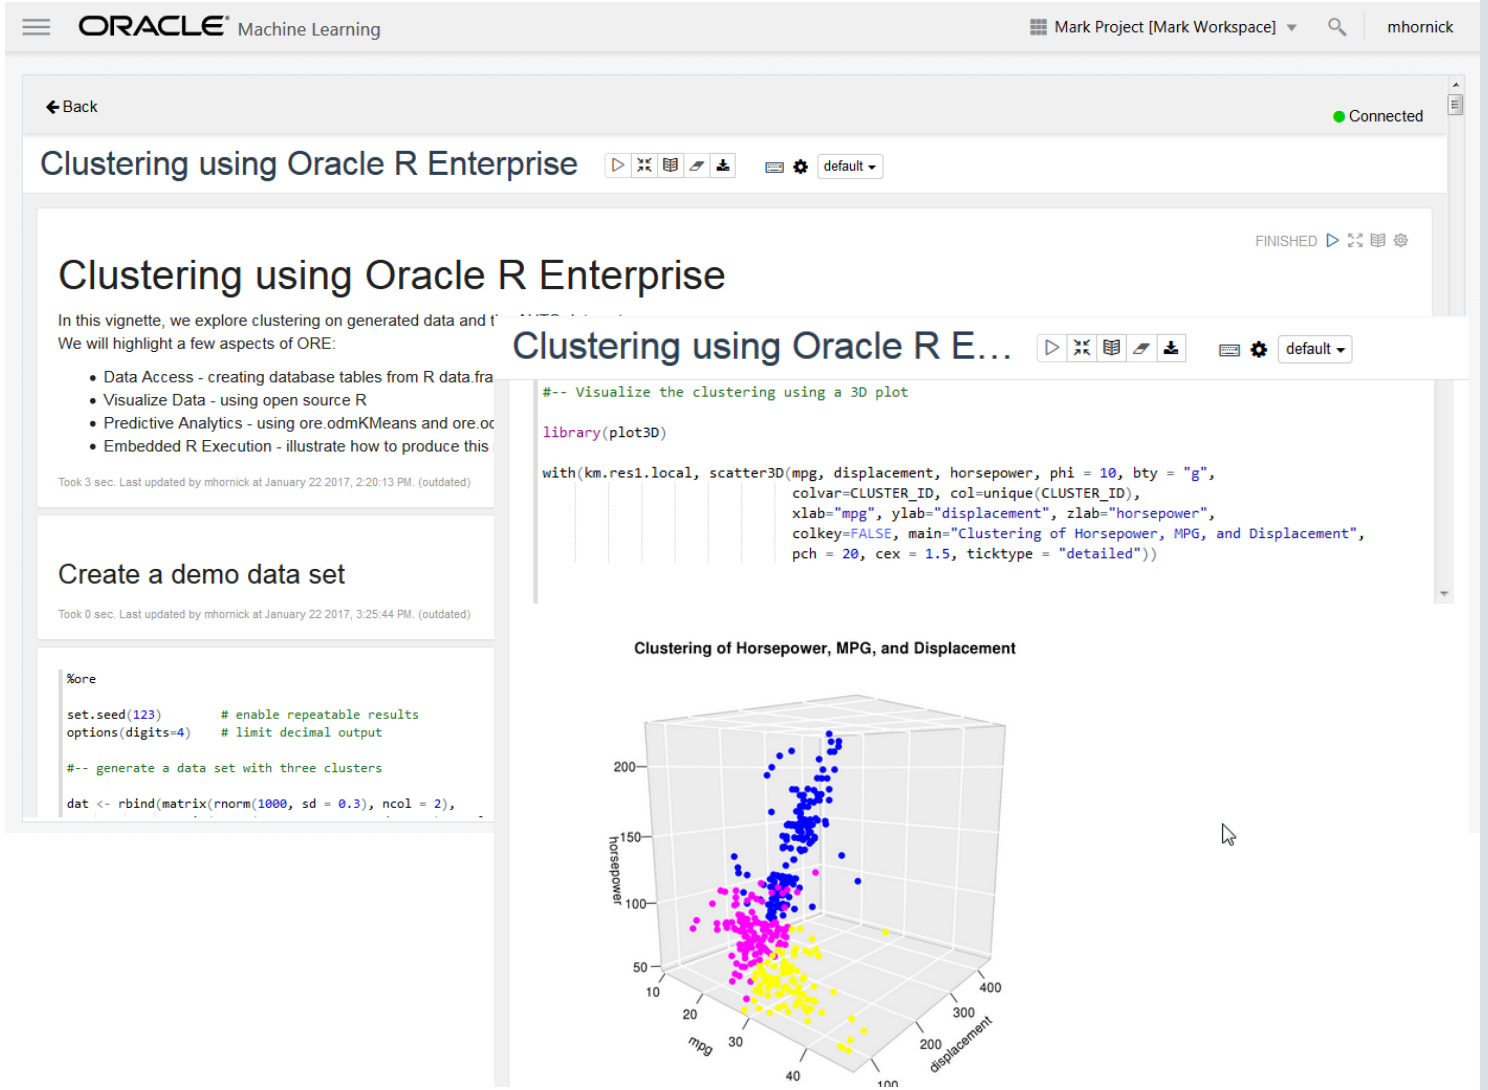

Oracle Machine Learning

Multi-Platform, Multi-Engine Oracle Machine Learning for the Cloud

Introduction

- A collaborative unified notebook user interface for Oracle's machine learning capabilities
- Enables teams to…
  - Leverage the languages, platforms and engines of their choice
  - Explore and prepare data
  - Build, access, and use machine learning models
  - Deploy machine learning solutions

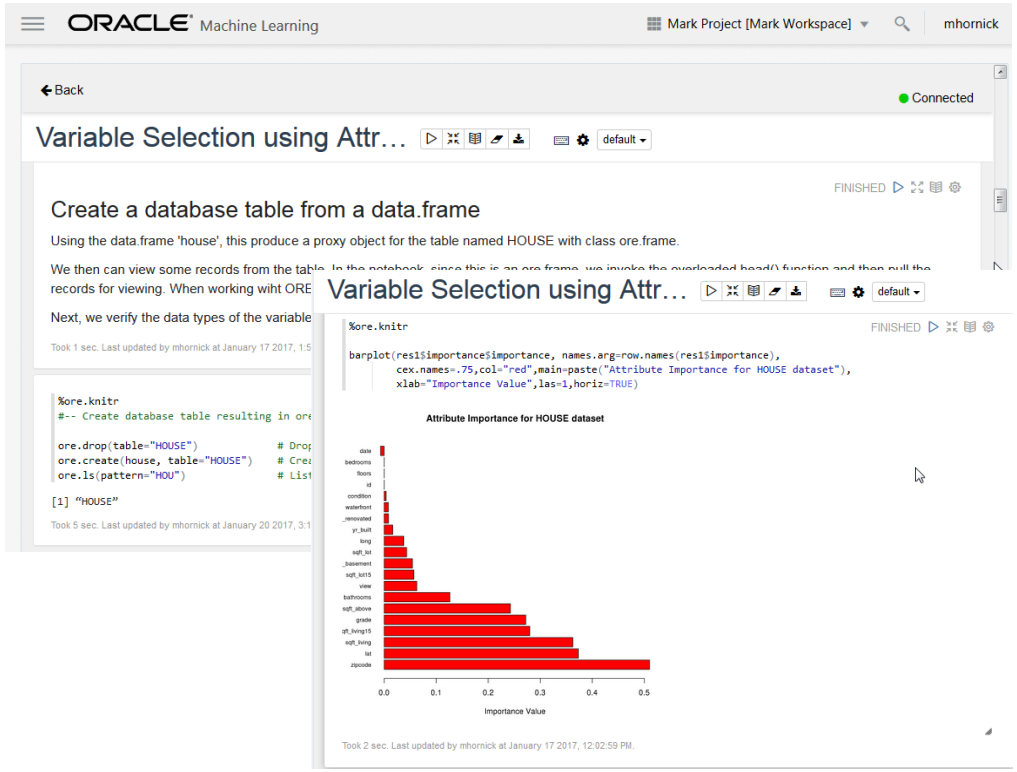

Oracle Machine Learning

Multi-Platform, Multi-Engine Oracle Machine Learning for the Cloud

Features

- Zeppelin-based Unifying Notebook UI
  - Multiple backend servers, analytical engines and languages
  - Oracle and open source algorithms, platforms, tools, and data sources (SQL, R, Spark, Python)
  - Provides a flexible toolbox for the data scientist
- Supports collaborative development
  - Shared notebooks and templates, with access permissions
- Supports deployment of predictive analytics solutions
  - Enables publishing libraries, templates, examples of common use cases
- Dynamic and expandable platform

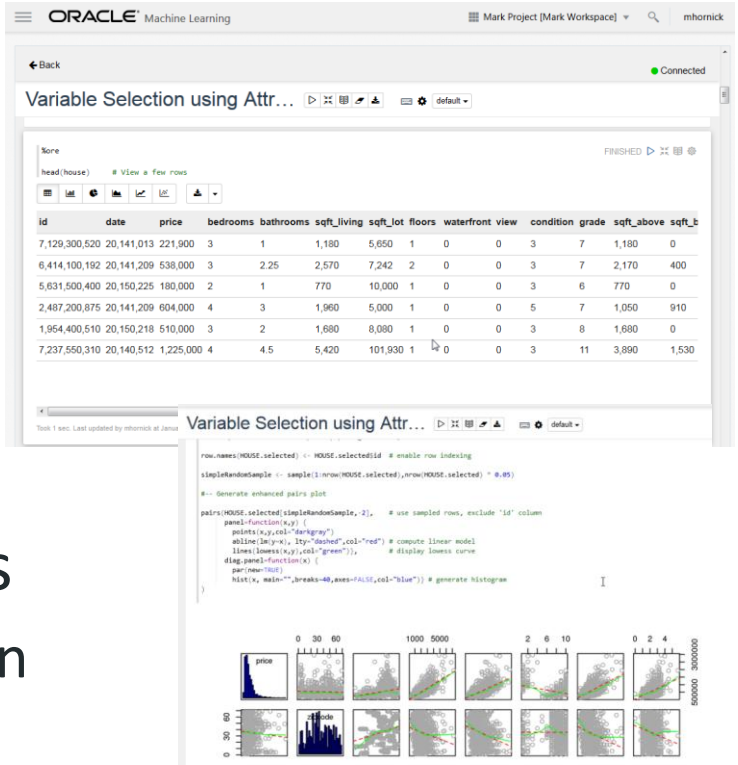

Oracle Machine Learning

Multi-Platform, Multi-Engine Oracle Machine Learning for the Cloud

Benefits

- Harnesses the power of Oracle's multi-platform data and machine learning ecosystem and open source machine learning algorithms
- Enables collaboration between data scientists, citizen data scientists, and application developers
- Develop and publish analytical methodologies, templates and micro-services in the Oracle Cloud

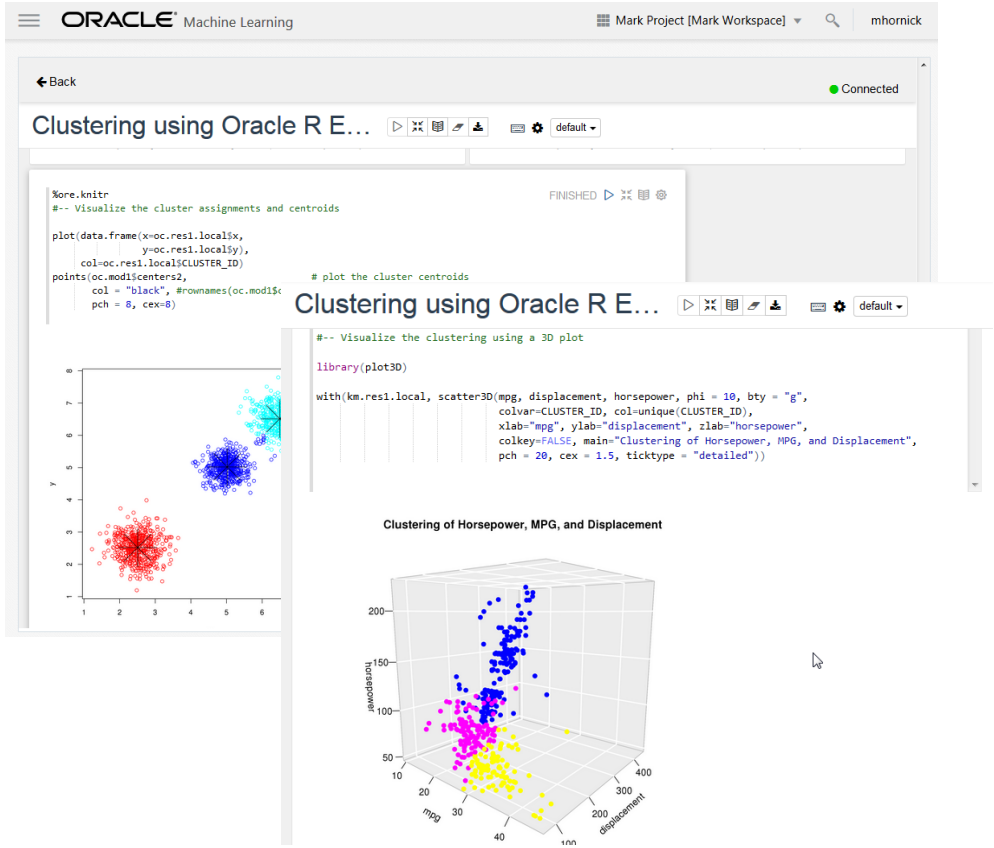

Oracle's Advanced Analytics and Machine Learning Platform

Multiple interfaces across platforms — SQL, R, GUI, Dashboards, Apps

Users

Information Producers

R programmers

R Client

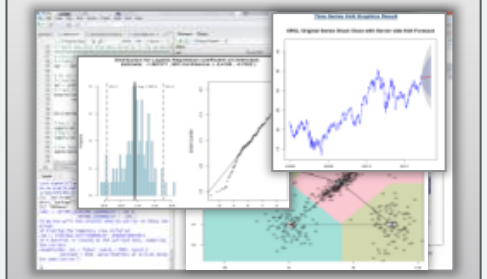

Data & Business Analysts

SQLDEV/
Oracle Data Miner

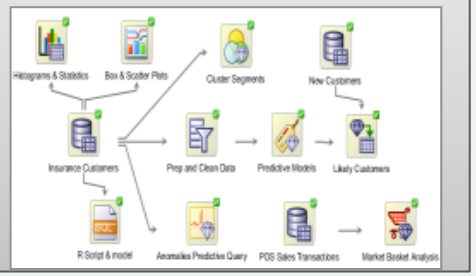

Data Scientists (R, SQL, Python, etc.)

Oracle
Machine Learning

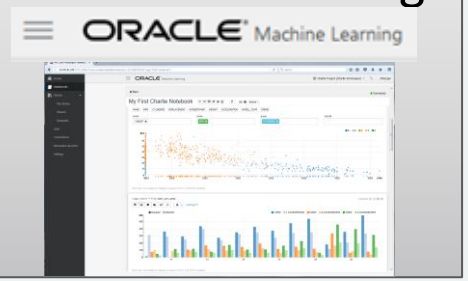

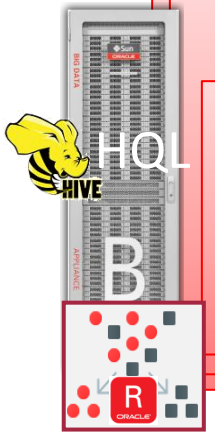

Hadoop

ORAAH
Parallel, distributed algorithms

Oracle Database Enterprise Edition

Oracle Advanced Analytics - Database Option
SQL Data Mining, ML & Analytic Functions + R Integration for Scalable, Distributed, Parallel in-DB ML Execution

Oracle Cloud

ORACLE

Oracle Machine Learning Strategy and Vision

One server side product, with a single analytic library, supporting multiple data platforms, analytical engines, UIs and deployment strategies

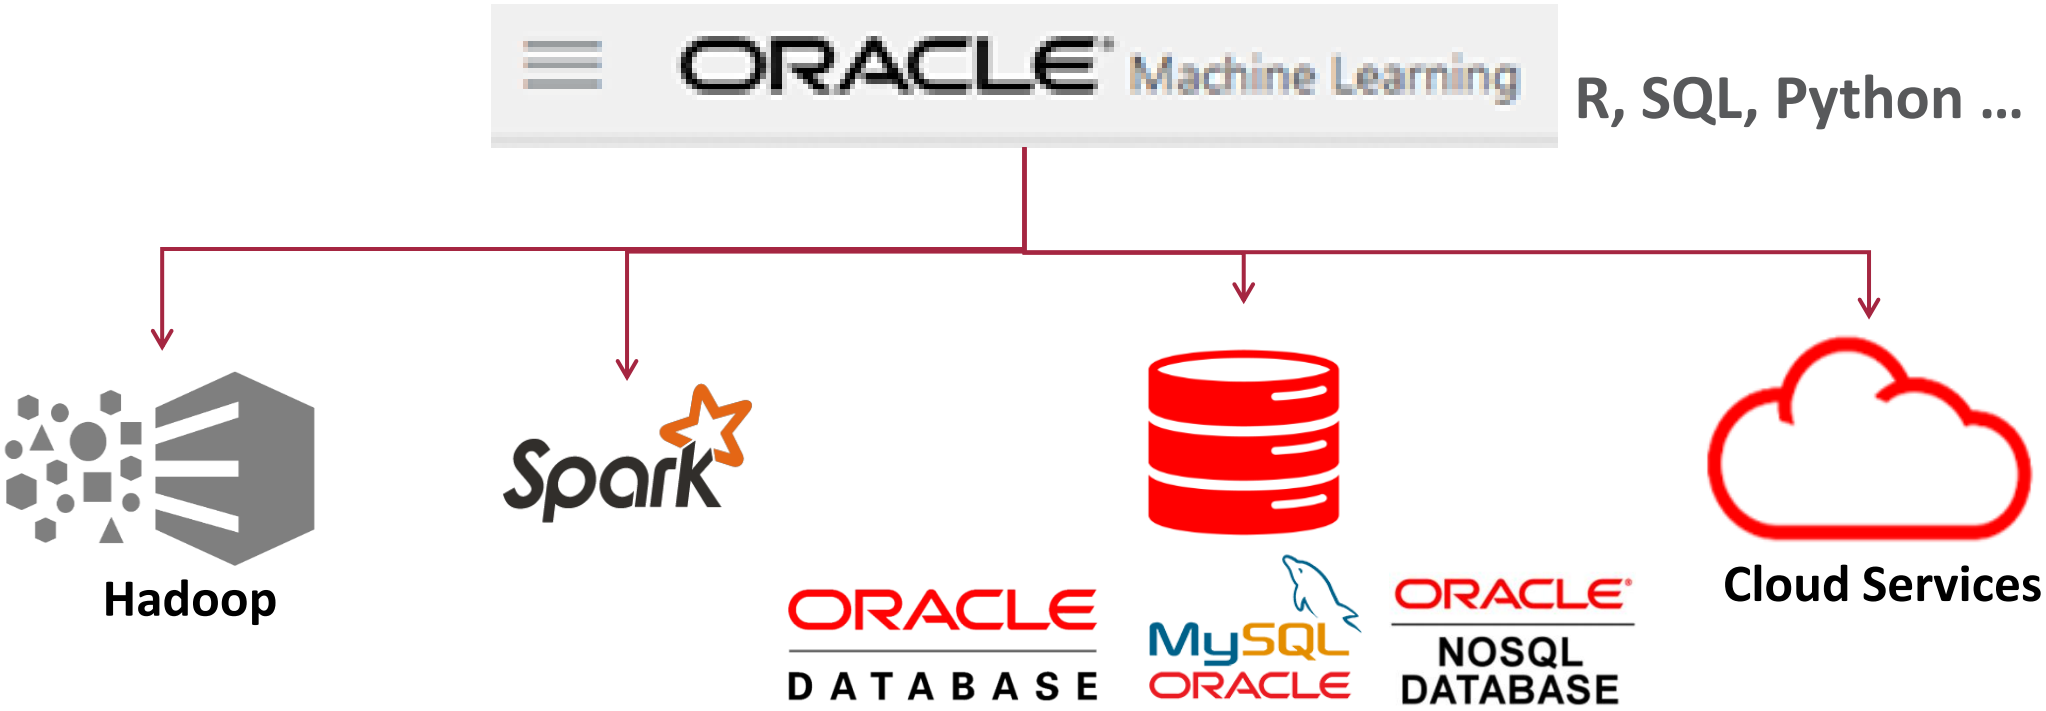

Oracle Machine Learning

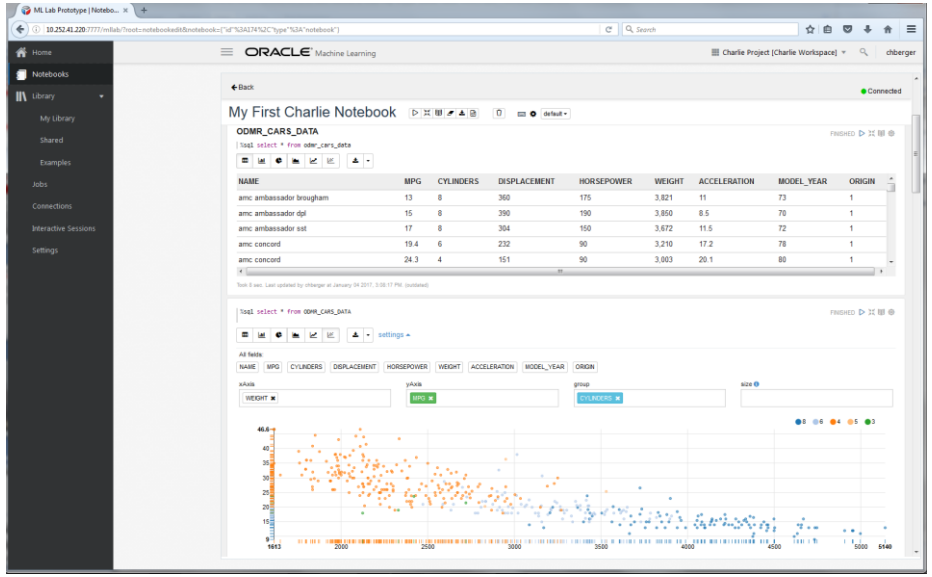

OML Beta Program 2017

- Seeking data scientists and data analysts for early product access, beta testing, and providing valuable customer feedback
- Oracle Machine Learning will host “i-Betas” in the Oracle Cloud for easier customer access and testing

OML Contacts in Product Management

- Charlie Berger – charlie.berger@oracle.com
- Mark Hornick – mark.hornick@oracle.com
- Marcos Arancibia –marcos.arancibia@oracle.com


ORACLE®