

# Oracle In-Database ML with OML4SQL/OML4PY in Autonomous DB



## EMEA Community Tour 2022



Kai Yu



ORACLE  
ACE Director

Distinguished Engineer  
Dell Technologies, Inc

# Agenda

- Machine Learning Enabling Business Applications
- Machine Learning in Oracle Database: OML4SQL and OML4Py
- Running Machine Learning with Oracle Autonomous Database
- An Example: Building Machine Learning Model with OML
- Summary/Q & A

# About Me: Kai Yu

**DELL**Technologies



- Distinguished Engineer, Dell Technical Leadership Community  
Dell Database Solutions Engineering
- 28+ years working in Tech Industry
- Specializing in Oracle Database, Virtualization/Cloud and Machine Learning
- Author and Frequent Speaker at IEEE and Oracle Conferences
- IOUG Cloud Computing SIG Co-founder and VP
- Oracle ACE Director since 2010
- Co-recipient of the OAUG Innovator of Year Award
- Oracle Technologist of the Year: Cloud Architect by Oracle Magazine
- My Blog: <http://kyuoracleblog.wordpress.com/>



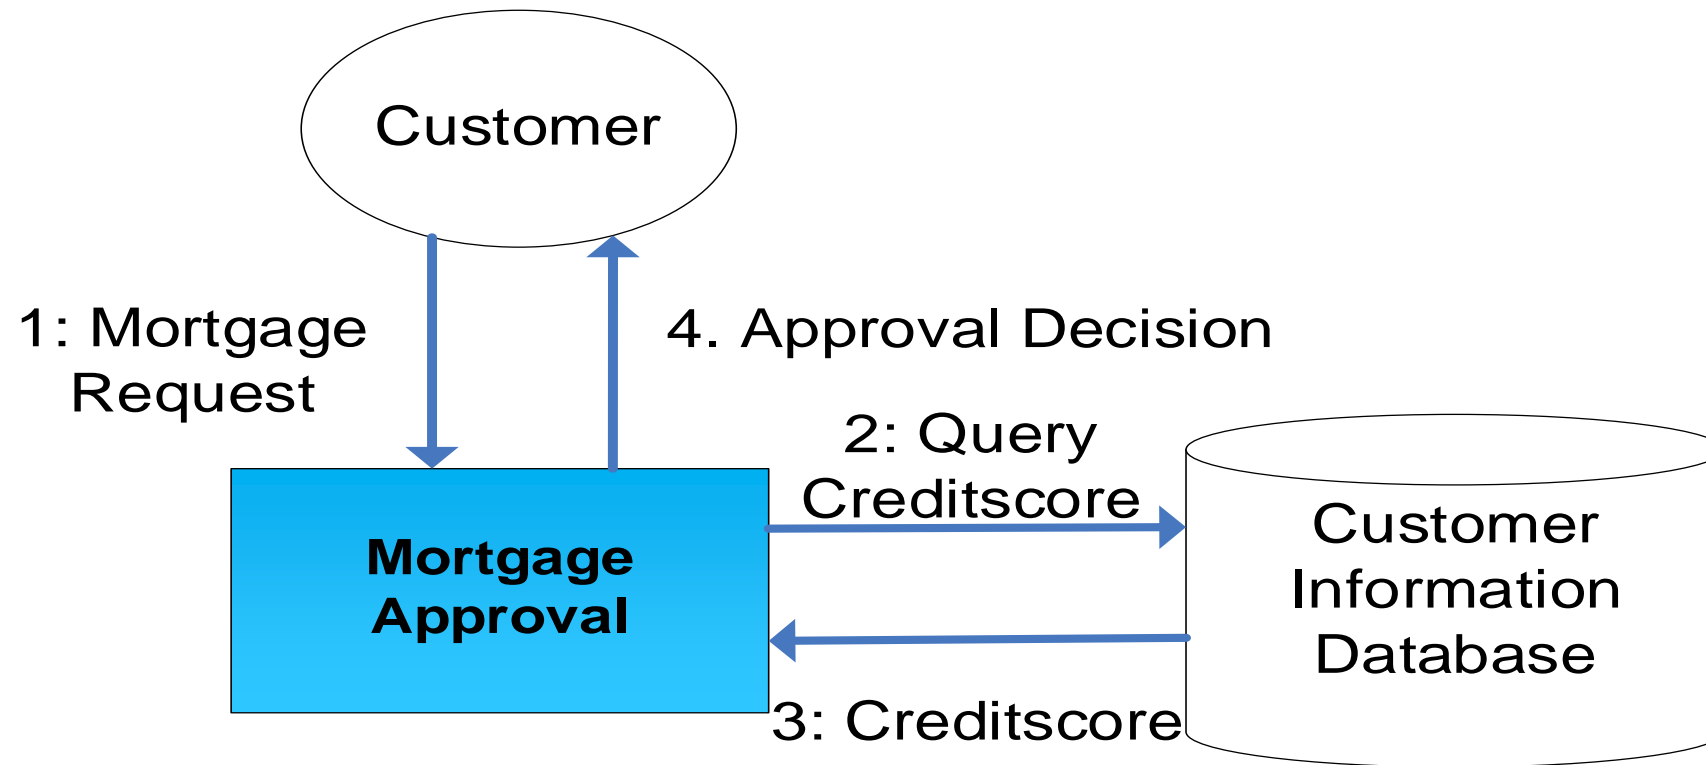
# Machine Learning Enabling Business Applications

# Machine Learning Enabling Business Applications

A simple mortgage request approval application

Mortgage approval decisions made based on the customer credit scores.

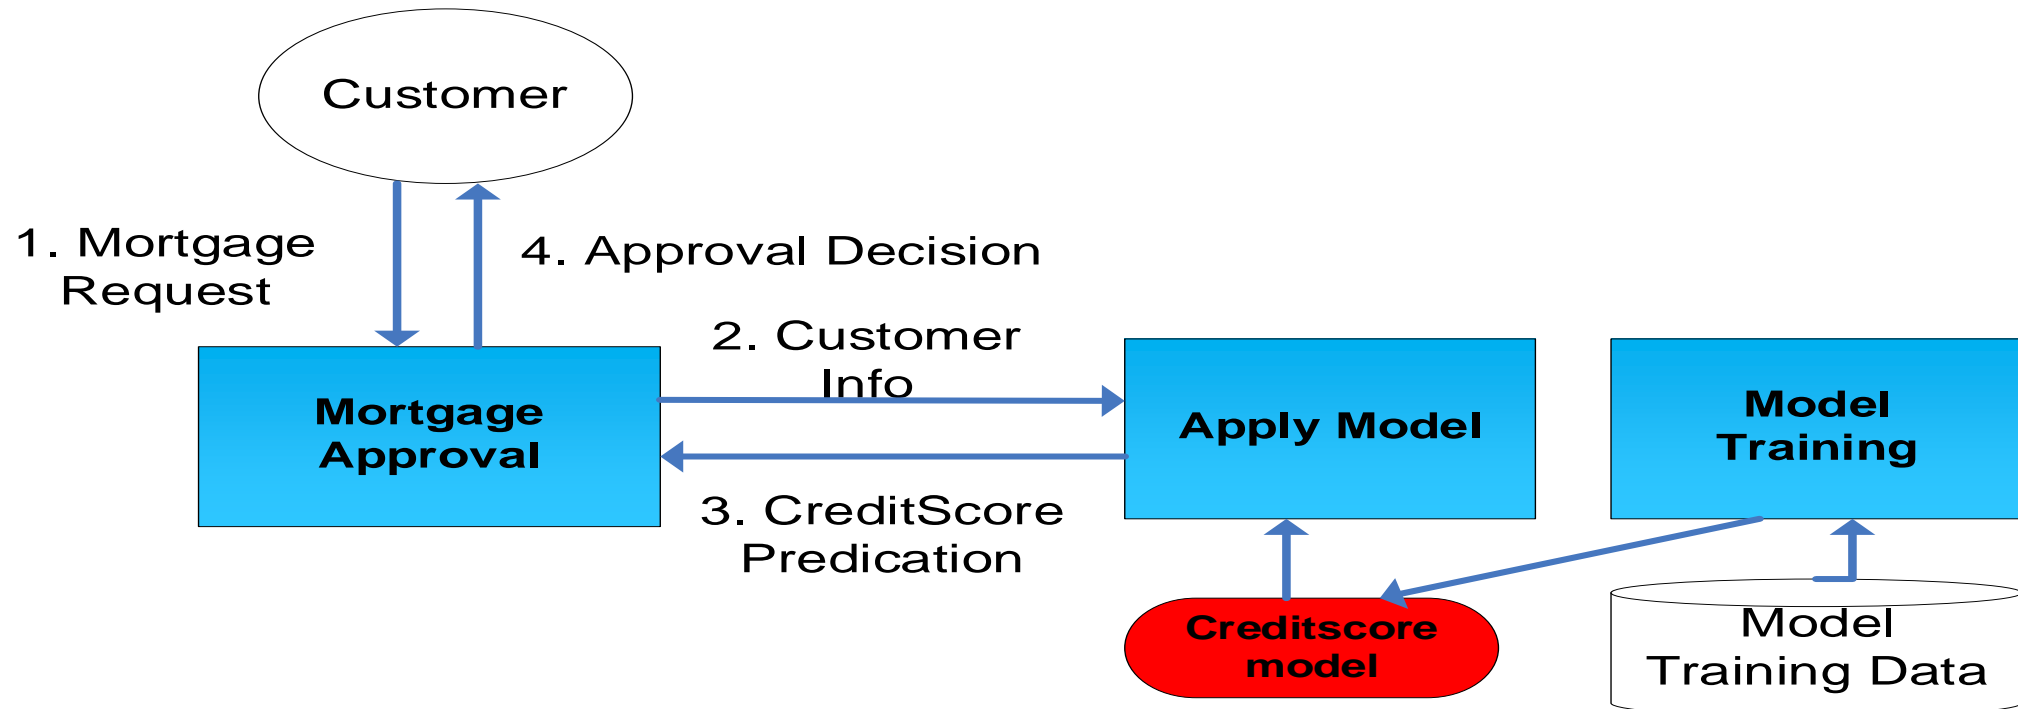
- A simple business process for customers with known credit scores.
- Limitation: What if the customer's credit score is unknown?



# Machine Learning Enabling Business Applications

An intelligent mortgage approval application for customers with unknown credit score

- Predict customer credit scores based on pre-trained credit score model.
- Add intelligence to the traditional business applications with predictive capability through Machine Learning.




# Machine Learning Enabling Business Applications

Next generation enterprise applications are more data driven and more intelligent.

- Transform massive volumes of data into meaningful insights
- Using AI and ML methods in business applications.

84% of organizations are using at least one emerging technology — AI, IoT, digital assistants, blockchain according to a recent survey by Oracle.

(<https://www.oracle.com/a/ocom/docs/esg-research-oracle-emerging-technologies.pdf>)



## Top emerging technology statistics: Finance

### Offload to the cloud:


**92%**  
**SaaS enablement:**  
92% of respondents see SaaS finance applications as enablers for the adoption of emerging technologies.

**2:1**  
**Prebuilt capabilities:**  
Organizations that add emerging technology capabilities to finance applications by purchasing prebuilt solutions outnumber those that build solutions by two to one.

### Take advantage of emerging technology benefits:

**37%**  
**AI effectiveness:**  
Users of AI within financial systems report an average improvement of 33% in productivity and 37% reduction in errors.

**4 DAYS**  
**AI efficiency:**  
Organizations have reduced the time needed to complete the monthly financial close process by about four days on average thanks to the incorporation of AI into their finance systems.



## Top emerging technology statistics: Supply chain management

### Offload to the cloud:

**88%**  
**SaaS enablement:**  
88% of respondents see SaaS SCM applications as enablers for the adoption of emerging technologies.

**3:1**  
**Prebuilt capabilities:**  
Organizations that add emerging technology capabilities to SCM applications by purchasing solutions outnumber those that build solutions by three to one.

### Specific impacts by emerging tech:

**76%**  
**AI efficiency:**  
76% of organizations cite increased employee productivity as a realized benefit of AI-powered SCM.

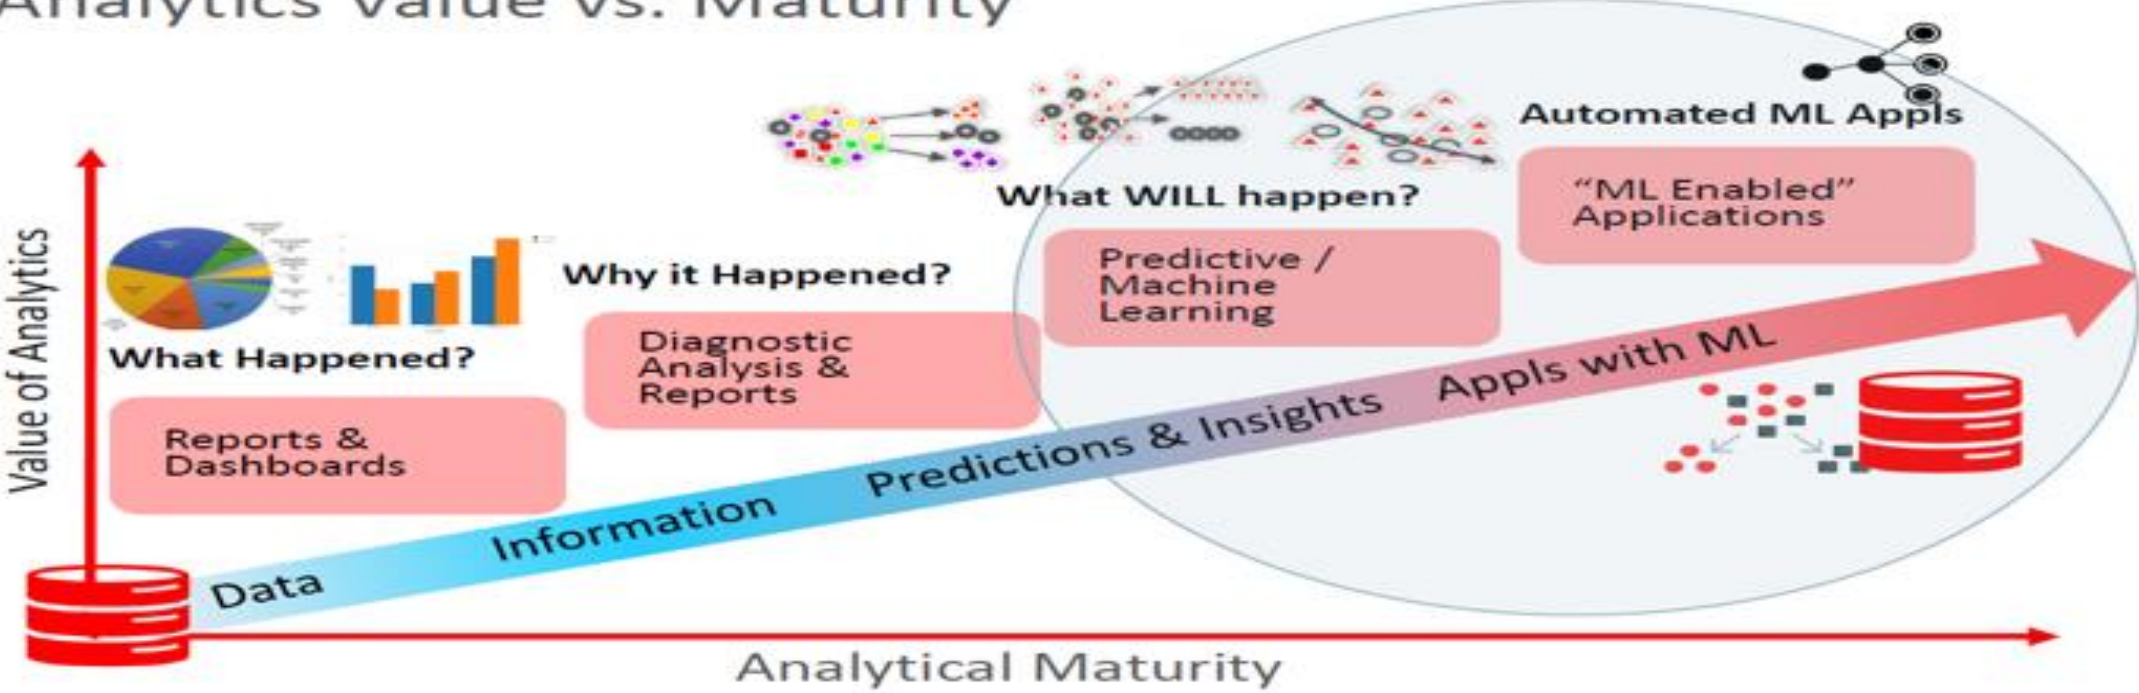
**75%**  
**Increase adoption:**  
75% of respondents credit chatbots and intelligent voice assistants with increasing the use of SCM apps by suppliers and customers.

# Machine Learning Enabling Business Applications

Advanced Analytics and Machine Learning bring great business values

- Transform massive volumes of data into meaningful insights:
  - Started with historical report and moved to advanced analytics such as Prediction, Pattern Discovery, Grouping and Actionable information
- ML enabled applications.

## Analytics Value vs. Maturity





# Machine Learning Enabling Business Applications

## Business Problems involving machine learning

- Pattern discovery: discovers certain patterns or trends from a large set of data
- Prediction: predict the probability of a future event based on historical data
- Grouping : find the nature group of data
- Actionable information: generate actionable decision from a large set of data

## Supervised Learning vs unsupervised learning

### Supervised learning:

- Directed by previously known output target and its dependent input attributes
- ML model: the target value as a function of a set of independent attributes based on a known data set .
- Supervised Learning Functions: Attribute Importance, Classification, Regression

# Machine Learning Enabling Business Applications

## Unsupervised Learning:

- Used when we do not know the target value.
- There is no previously known results like the ones in supervised learning to train the model and test the model.
- Unsupervised learning functions:
  - Anomaly Detection: for example, finding customer's different purchasing pattern
  - Association Rules: discover the rules that defines the items that tend to occur together, for example, the products that customers tend to buy together.
  - Clustering: groups the items
  - Feature Extraction: creates new attributions using liner combinations of the original attributions

# Machine Learning Enabling Business Applications

Some special challenges and considerations:

- Enterprise data may be generated by transactional applications
- Large Data set used for model training and model validation
- Data are stored in structured format like in relational database tables.
- Data preparation and data exploratory on large data set
- Real time data may be needed
- The ML models may be deployed with traditional applications

# Machine Learning in Oracle Database

## OML4SQL and OML4Py

# Machine Learning in Oracle Database

Move the Algorithm, Not the Data → In Database Machine Learning

Oracle “moves the algorithms to the data” : Make the machine learning algorithms available within Oracle Database

Minimizes or eliminates data movement

## Machine Learning Algorithms Require Data

Move the Algorithms, Not the Data!



$$P(c | x) = \frac{P(x | c)P(c)}{P(x)}$$

Likelihood:  $P(x | c)$   
Class Prior Probability:  $P(c)$   
Posterior Probability:  $P(c | x)$   
Predictor Prior Probability:  $P(x)$

$$P(c | X) = P(x_1 | c) \times P(x_2 | c) \times \dots \times P(x_n | c) \times P(c)$$

An “AI Database” or “Thinking Database”? → It Changes Everything!

# Machine Learning in Oracle Database

Built-in machine learning functions in database process large data sets

- Built-in functions in Oracle database support supervised and unsupervised ML functions such as classification, clustering, prediction, regression etc.
- Provide over 30 ML algorithms: Decision tree, k-means, linear regression, logistic regression,
- Data processing tool SQL interface for data gathering and preparation
- Support entire major machine learning process: Data Gathering and Preparation, Model Training, Model Evaluation, Model Scoring and deployment with SQL and/or database built-in functions & procedures.

# Machine Learning in Oracle Database

Machine Learning: algorithms automatically examine large amount of data to identify patterns, discover new insight and make predictions

Machine learning algorithms implemented as SQL functions inside Oracle DB

## CLASSIFICATION

- Naïve Bayes
- Logistic Regression (GLM)
- Decision Tree
- Random Forest
- Neural Network
- Support Vector Machine
- Explicit Semantic Analysis

## CLUSTERING

- Hierarchical K-Means
- Hierarchical O-Cluster
- Expectation Maximization (EM)

## ANOMALY DETECTION

- One-Class SVM

## TIME SERIES

- Forecasting - Exponential Smoothing
- Includes popular models  
e.g. Holt-Winters with trends, seasonality, irregularity, missing data

## REGRESSION

- Linear Model
- Generalized Linear Model
- Support Vector Machine (SVM)
- Stepwise Linear regression
- Neural Network
- LASSO

## ATTRIBUTE IMPORTANCE

- Minimum Description Length
- Principal Comp Analysis (PCA)
- Unsupervised Pair-wise KL Div
- CUR decomposition for row & AI

## ASSOCIATION RULES

- A priori/ market basket

## PREDICTIVE QUERIES

- Predict, cluster, detect, features

## SQL ANALYTICS

- SQL Windows
- SQL Patterns
- SQL Aggregates

## FEATURE EXTRACTION

- Principal Comp Analysis (PCA)
- Non-negative Matrix Factorization
- Singular Value Decomposition (SVD)
- Explicit Semantic Analysis (ESA)

## TEXT MINING SUPPORT

- Algorithms support text
- Tokenization and theme extraction
- Explicit Semantic Analysis (ESA) for document similarity

## STATISTICAL FUNCTIONS

- Basic statistics: min, max, median, stdev, t-test, F-test, Pearson's, Chi-Sq, ANOVA, etc.

## R AND PYTHON PACKAGES

- Third-party R and Python Packages through Embedded Execution
- Spark MLlib algorithm integration

# Oracle Machine Learning for SQL (OML4SQL)

## Simple SQL Syntax – Classification Model:

- Setting mode parameter values with insert statements      Model build(PL/SQL):

```
/* Create a Build Setting (DT) for Model Build */  
  
EXECUTE IMMEDIATE 'CREATE TABLE n1_build_settings (setting_name VARCHAR2(30),setting_value VARCHAR2(4000))';  
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings (setting_name, setting_value) VALUES ('ALGO_NAME', 'ALGO_DECISION_TREE')';  
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings (setting_name, setting_value) VALUES ('PREP_AUTO', 'ON')';  
  
DBMS_OUTPUT.PUT_LINE ('Created model build settings table: n1_build_settings ');  
  
/*  
-- Populate and Adjust Model Setting (DT) for Model Build  
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings VALUES ('TREE_TERM_MAX_DEPTH', 7)';  
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings VALUES ('TREE_TERM_MINREC_SPLIT', 20)';  
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings VALUES ('TREE_TERM_MINPCT_SPLIT', .1)';  
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings VALUES ('TREE_TERM_MINREC_NODE', 10)';  
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings VALUES ('TREE_TERM_MINPCT_NODE', 0.05)';  
*/
```

## BEGIN

```
DBMS_DATA_MINING.CREATE_MODEL(  
  model_name           =>'N1_CLASS_MODEL',  
  mining_function      =>'dbms_data_mining.classification',  
  data_table_name     =>'N1_TRAIN_DATA',  
  case_id_column_name =>'CUSTOMER_ID',  
  target_column_name  =>'CREDIT_SCORE_BIN',  
  setting_table_name  =>'N1_BUILD_SETTINGS')
```

END;

- Real-time scoring(SQL query):

```
select prediction_probability(N1_CLASS_MODEL, 'Good Credit'  
  USING 'Rich' as WEALTH, 2000 as income, 'Silver' as customer_value_segment) Prediction_Probability  
from dual;
```



# Oracle Machine Learning for SQL (OML4SQL)

## In-database, parallel, distributed algorithms

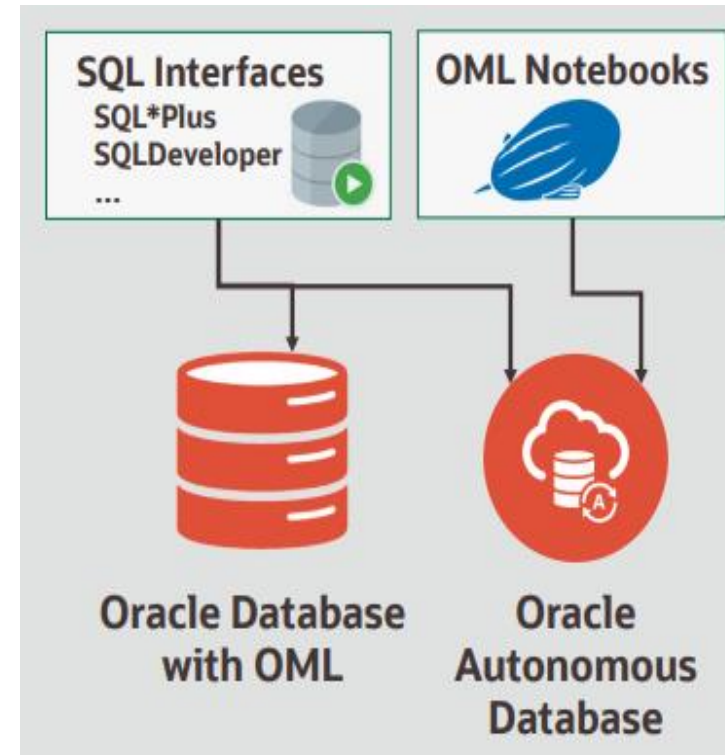
- No extracting data to separate ML engine
- Fast and scalable
- Batch and real-time scoring
- Explanatory prediction details

## ML models as first class database objects

- Access control via permissions
- Audit user actions
- Export / import models across databases

## New OML4SQL features added in 21c

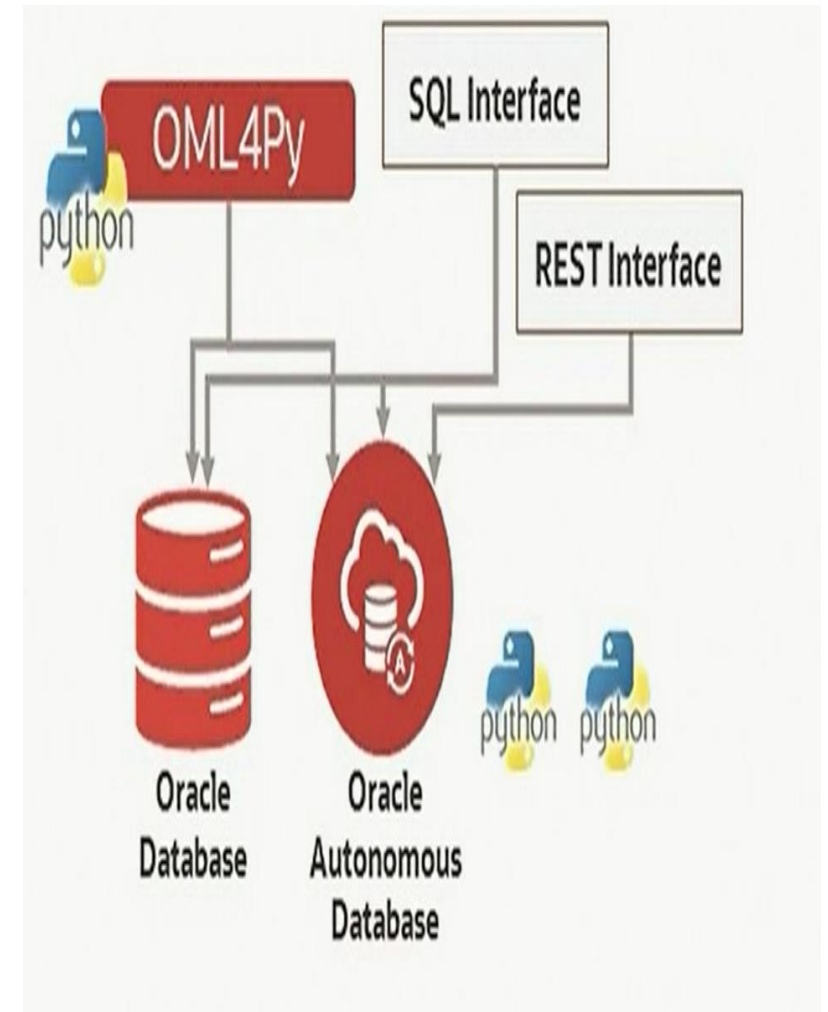
- OML MSET-SPRT Algorithm for Sequential Probability Ratio Test ;
- OML XGBoost Algorithm: a highly-efficient, scalable gradient tree boosting machine learning algorithm for regression and classification.
- Adam Optimization Solver for the Neural Network Algorithm; new Neural Network Algorithm Settings



# Oracle Machine Learning for Python (OML4Py)

## Oracle Machine Learning for Python (OML4Py)

- Enables the opensource Python environment to operate on database data at scale.
- Allow you to run Python commands and scripts for statistical analysis and machine learning on data stored in Oracle Database.
- Use in-database machine learning algorithms with OML4Py
- Leverage database server and minimizing data movement
- Use AutoML for automatic algorithm selection, feature selection, and model tuning
- Execute user-defined Python functions in non-parallel, data-parallel, and task-parallel fashion
- Store and manage Python objects in the database



# Oracle Machine Learning for Python (OML4Py)

OML4Py OML Classes that provide access to In-Database Machine Learning Algorithm, example:

1 .Create an OML SVM model object

```
# Create an oml SVM model object.
svm_mod = oml.svm('classification',
                 svms_kernel_function =
                 'dbms_data_mining.svms_linear')
```

2. Train the model with training data

```
# Fit the SVM model with the training data.
svm_mod = svm_mod.fit(train_x, train_y, case_id = 'ID')
```

3. Export the model to svm\_model table

```
# Export the oml.svm model to a new table named 'svm_sermod'
# in the database.
svm_export = svm_mod.export_sermodel(table='svm_sermod')
type(svm_export)
```

4. Import the model

```
BEGIN
  SELECT SERVAL INTO v_blob FROM "svm_sermod";
  dbms_data_mining.import_sermodel(v_blob, 'my_iris_svm_classifier');
END;
```

5. Apply the model to make predictions on new data

```
-- Make predictions and display cases where mod(ID,3) equals 0.
SELECT ID, "Species" AS TARGET_SPECIES,
       PREDICTION(my_iris_svm_classifier USING "Sepal_Length",
                 "Sepal_Width",
                 "Petal_Length", "Petal_Width")
       AS PREDICT_SPECIES
FROM "IRIS_TEST_DATA" WHERE MOD(ID,3) = 0;
```

# Oracle Machine Learning for Python (OML4Py)

OML4Py OML classes that provide access to In-Database Machine Learning Algorithm:

<b>Class</b>	<b>Algorithm</b>	<b>Function of Algorithm</b>	<b>Description</b>
<code>oml.ai</code>	Minimum Description Length	Attribute importance for classification or regression	Ranks attributes according to their importance in predicting a target.
<code>oml.ar</code>	Apriori	Association rules	Performs market basket analysis by identifying co-occurring items (frequent itemsets) within a set.
<code>oml.dt</code>	Decision Tree	Classification	Extracts predictive information in the form of human-understandable rules. The rules are if-then-else expressions; they explain the decisions that lead to the prediction.
<code>oml.em</code>	Expectation Maximization	Clustering	Performs probabilistic clustering based on a density estimation algorithm.
<code>oml.esa</code>	Explicit Semantic Analysis	Feature extraction	Extracts text-based features from a corpus of documents. Performs document similarity comparisons.
<code>oml.glm</code>	Generalized Linear Model	Classification Regression	Implements logistic regression for classification of binary targets and linear regression for continuous targets.
<code>oml.km</code>	<i>k</i> -Means	Clustering	Uses unsupervised learning to group data based on similarity into a predetermined number of clusters.
<code>oml.nb</code>	Naive Bayes	Classification	Makes predictions by deriving the probability of a prediction from the underlying evidence, as observed in the data.
<code>oml.nn</code>	Neural Network	Classification Regression	Learns from examples and tunes the weights of the connections among the neurons during the learning process.
<code>oml.rf</code>	Random Forest	Classification	Provides an ensemble learning technique for classification of data.
<code>oml.svd</code>	Singular Value Decomposition	Feature extraction	Performs orthogonal linear transformations that capture the underlying variance of the data by decomposing a rectangular matrix into three matrices.
<code>oml.svm</code>	Support Vector Machine	Anomaly detection Classification Regression	Builds a model that is a profile of a class, which, when the model is applied, identifies cases that are somehow different from that profile.

# Machine Learning in Oracle Database

## Use popular languages

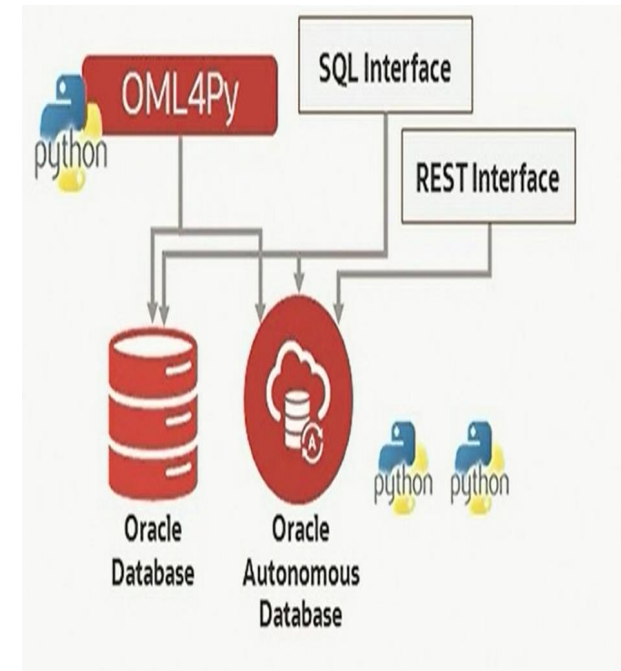
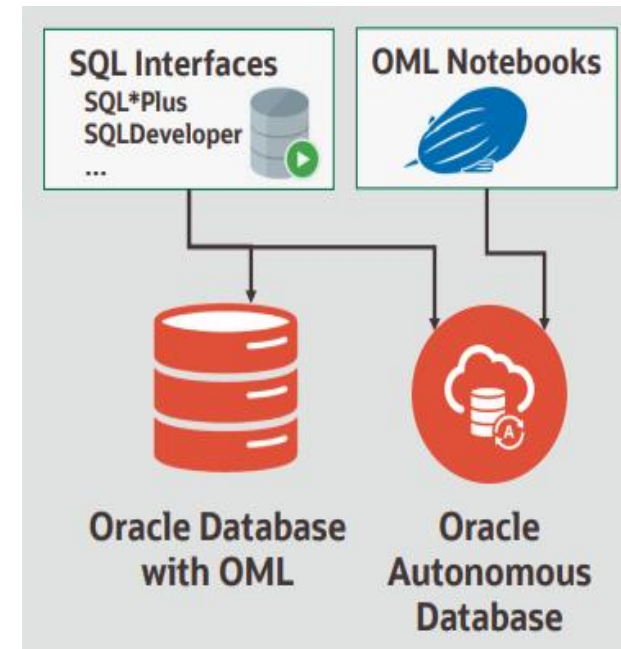
- SQL: Oracle Machine Learning for SQL: (OML4SQL)
- Python: Oracle Machine Learning for Python (OML4Py)
- R: Oracle Machine Learning for R (OML4R)
- REST :API interface

## GUI Interfaces

- Oracle Machine Learning Notebooks
- Oracle Data Miner User Interface (Oracle SQL Developer extension)
- Oracle Machine Learning AutoML User Interface

## Support entire major machine learning process:

- Data Gathering and Preparation,
- Model Training, Model Evaluation,
- Model Scoring and deployment with SQL and/or database built-in functions & procedures.



# Automated Machine Learning (AutoML)

## Automated Machine Learning in OLM4Py

- . Introduced in Oracle 21c
- . Provided the following AutoML capabilities:
  - Automate algorithm selection,
  - Automate feature selection
  - Automated hyperparameter

## Machine Learning Algorithms Supported by AutoML \*\*

Algorithm Abbreviation	Algorithm Name
dt	Decision Tree
glm	Generalized Linear Model
glm_ridge	Generalized Linear Model with ridge regression
nb	Naive Bayes
nn	Neural Network
rf	Random Forest
svm_gaussian	Support Vector Machine with Gaussian kernel
svm_linear	Support Vector Machine with linear kernel

## The Automated Machine Learning classes \*\*

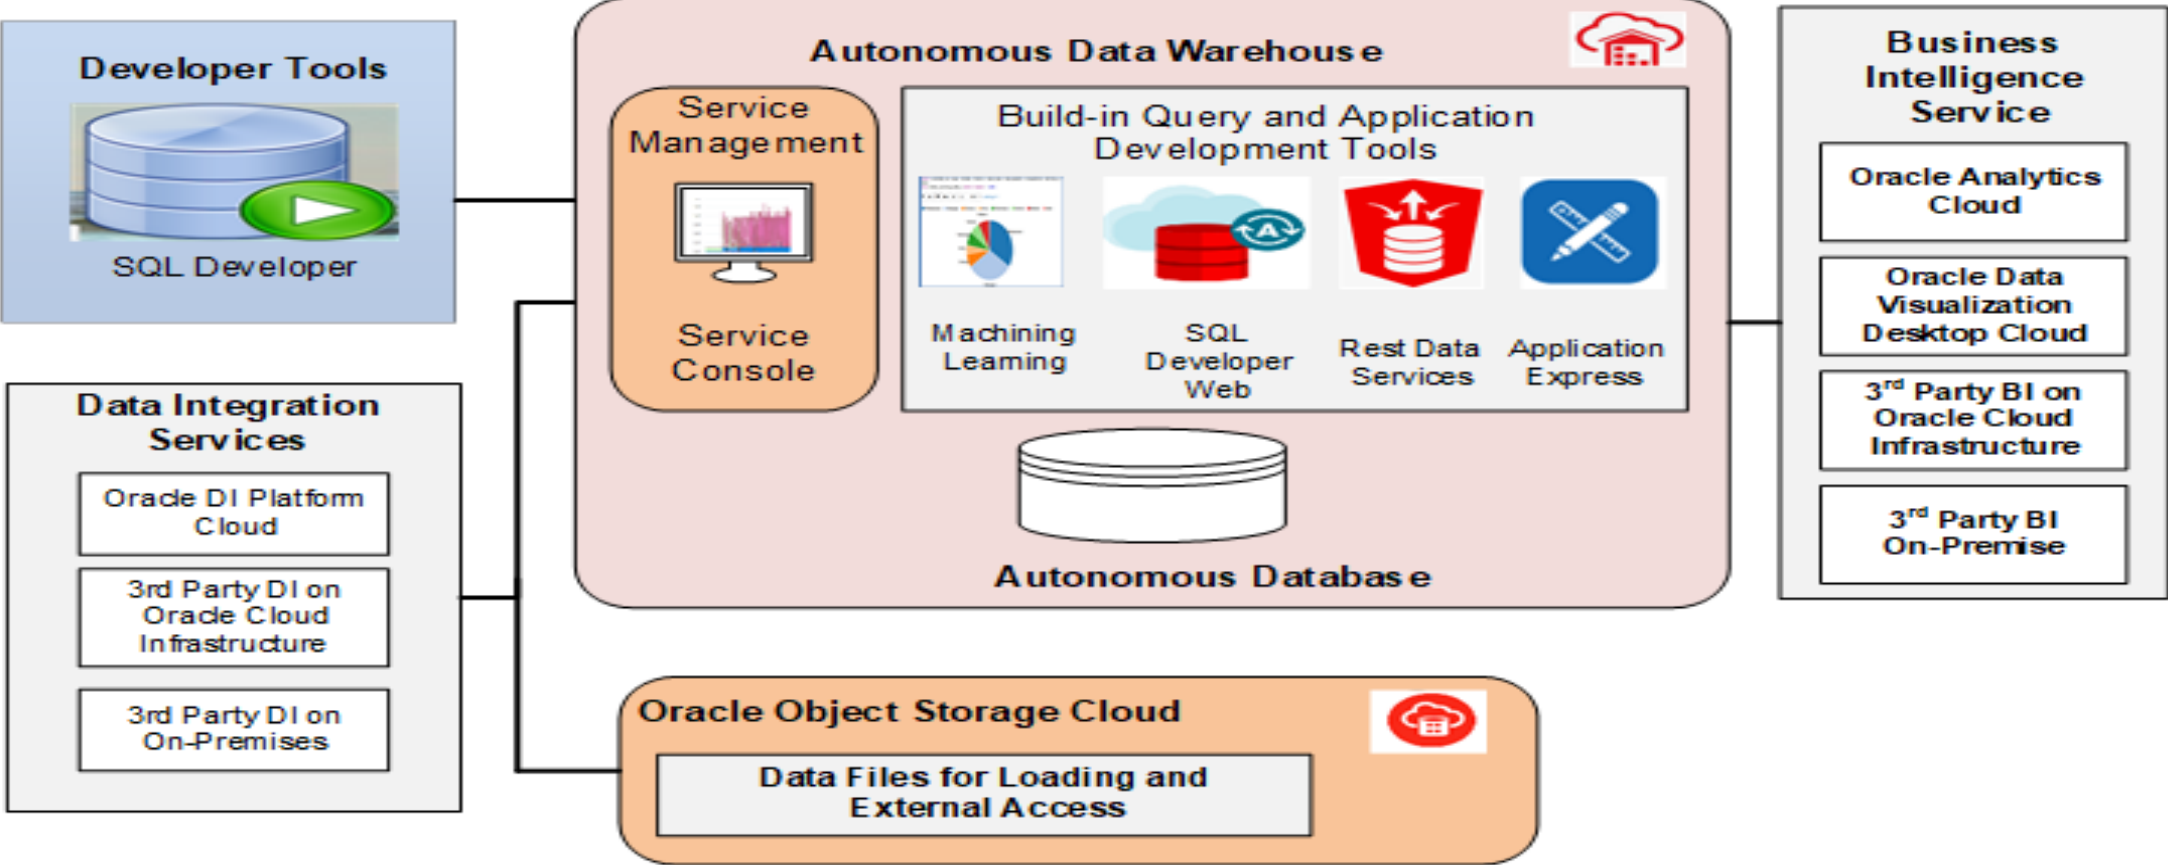
Class	Description
<code>oml.automl.AlgorithmSelection</code>	Using only the characteristics of the data set and the task, automatically selects the best algorithms from the set of supported Oracle Machine Learning algorithms. Supports classification and regression functions.
<code>oml.automl.FeatureSelection</code>	Uses meta-learning to quickly identify the most relevant feature subsets given a training data set and an Oracle Machine Learning algorithm. Supports classification and regression functions.
<code>oml.automl.ModelTuning</code>	Uses a highly parallel, asynchronous gradient-based hyperparameter optimization algorithm to tune the algorithm hyperparameters. Supports classification and regression functions.
<code>oml.automl.ModelSelection</code>	Selects the best Oracle Machine Learning algorithm and then tunes that algorithm. Supports classification and regression functions.

\*\* Refer to Oracle® Machine Learning for Python User's Guide Release 1.0 E97014-16 December 2021

# Running Machine Learning with Oracle Autonomous Databases

# OML in Oracle Autonomous Databases Overview

Oracle Autonomous Database Environment:





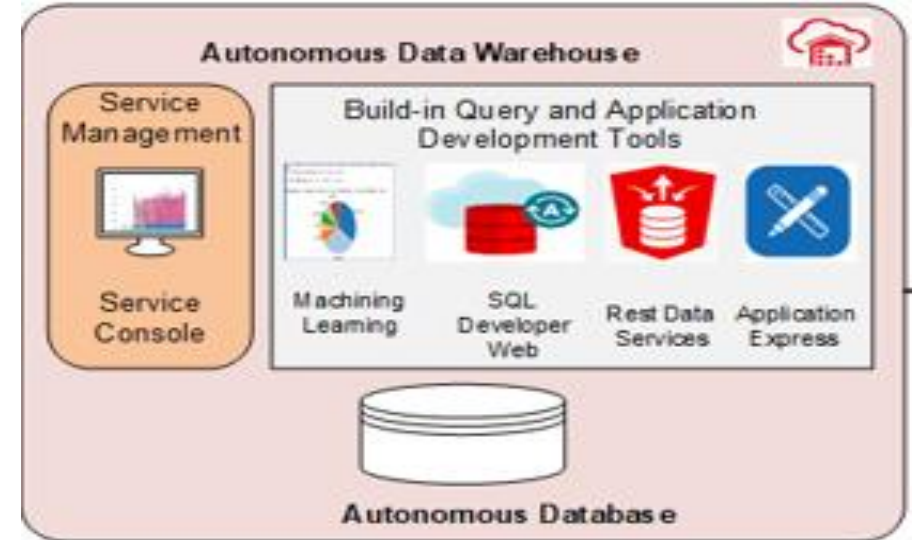
# Running Oracle Machine Learning with Autonomous Database

Oracle Machine Learning: a part of Autonomous Database collaborate environment

- A web-based development to create data mining notebook
- Used by data scientists, developer and business to users to perform data analytics, data discover and data virtualizations.
- Leverage the ADW scalability and performance in Oracle cloud

## Key Features:

- Collaborative UI for data scientists
- Packaged with Autonomous Data Warehouse Cloud Easy access to shared notebooks, templates, permissions, scheduler, etc.
- SQL ML algorithms API



# Running Oracle Machine Learning in Autonomous Database

Machine Learning: algorithms automatically examine large amount of data to identify patterns, discover new insight and make predictions

Machine learning algorithms implemented as SQL functions inside Oracle DB

## CLASSIFICATION

- Naïve Bayes
- Logistic Regression (GLM)
- Decision Tree
- Random Forest
- Neural Network
- Support Vector Machine
- Explicit Semantic Analysis

## CLUSTERING

- Hierarchical K-Means
- Hierarchical O-Cluster
- Expectation Maximization (EM)

## ANOMALY DETECTION

- One-Class SVM

## TIME SERIES

- Forecasting - Exponential Smoothing
- Includes popular models e.g. Holt-Winters with trends, seasonality, irregularity, missing data

## REGRESSION

- Linear Model
- Generalized Linear Model
- Support Vector Machine (SVM)
- Stepwise Linear regression
- Neural Network
- LASSO

## ATTRIBUTE IMPORTANCE

- Minimum Description Length
- Principal Comp Analysis (PCA)
- Unsupervised Pair-wise KL Div
- CUR decomposition for row & AI

## ASSOCIATION RULES

- A priori/ market basket

## PREDICTIVE QUERIES

- Predict, cluster, detect, features

## SQL ANALYTICS

- SQL Windows
- SQL Patterns
- SQL Aggregates

## FEATURE EXTRACTION

- Principal Comp Analysis (PCA)
- Non-negative Matrix Factorization
- Singular Value Decomposition (SVD)
- Explicit Semantic Analysis (ESA)

## TEXT MINING SUPPORT

- Algorithms support text
- Tokenization and theme extraction
- Explicit Semantic Analysis (ESA) for document similarity

## STATISTICAL FUNCTIONS

- Basic statistics: min, max, median, stdev, t-test, F-test, Pearson's, Chi-Sq, ANOVA, etc.

## R AND PYTHON PACKAGES

- Third-party R and Python Packages through Embedded Execution
- Spark MLlib algorithm integration

# Running Oracle Machine Learning with Autonomous Database

## Use Oracle Machine Learning: start with Oracle ML users from Oracle ML Admin

Overview > Autonomous Database > Autonomous Database Details

oml1 Always Free

Database Actions DB Connection Performance Hub Manage Scaling More actions

Autonomous Database Information Tools Tags

Database administration and developer tools for Autonomous Database

**Oracle APEX**  
Oracle APEX is a low-code development platform that you can use to build scalable, secure enterprise applications that can be deployed anywhere. [Learn more](#)  
[Open APEX](#)

**SODA Drivers**  
Simple Oracle Document Access (SODA) is a set of APIs that let you work with JSON documents managed by the Oracle Database without needing to use SQL. SODA drivers are available for REST, Java, Node.js, Python, PL/SQL, and C. [Learn more](#)  
[Download SODA Drivers](#)

**Oracle ML User Administration**  
Oracle Machine Learning is a development environment that uses a web-based interface to enable you to perform data analytics, data discovery and data visualizations. [Learn more](#)  
[Open Oracle ML User Administration](#)

**Graph Studio**  
Oracle Graph Studio lets you create property graph databases and automates the creation of graph models and in-memory graphs from database tables. To access Graph Studio, you must log in as a graph-enabled database user. Create this user in Database Actions. [Learn more](#)  
[Open Graph Studio](#)

Resources

Metrics [View all database metrics](#) [View audit and logs](#)

Start time: Sep 25, 2022 20:35:14 UTC | End time: Sep 25, 2022 21:35:14 UTC | Quick selects: Last hour | [Reset charts](#)

ORACLE Machine Learning User Administration

### Create User

[Create](#) [Cancel](#)

\* Username:

First Name:

Last Name:

\* Email Address:

Generate password and email account details to user. User will be required to reset the password on first sign in.

\* Password:

\* Confirm Password:

ORACLE Machine Learning User Administration

User Created

## Users

[+ Create](#) [X Delete](#)  Show All Users

User Name	Full Name	Role	Email	Created On	Status
ADMIN		System Administrator		11/10/20, 7:27 AM	Open
KAI_ML	Kai Yu	Developer	kai_yu@dell.com	9/25/22, 4:46 PM	Open

# Running Oracle Machine Learning with Autonomous Database

Use Oracle Machine Learning: Login to Oracle ML, taken into OML home page



ORACLE Cloud Infrastructure

**SIGN IN**

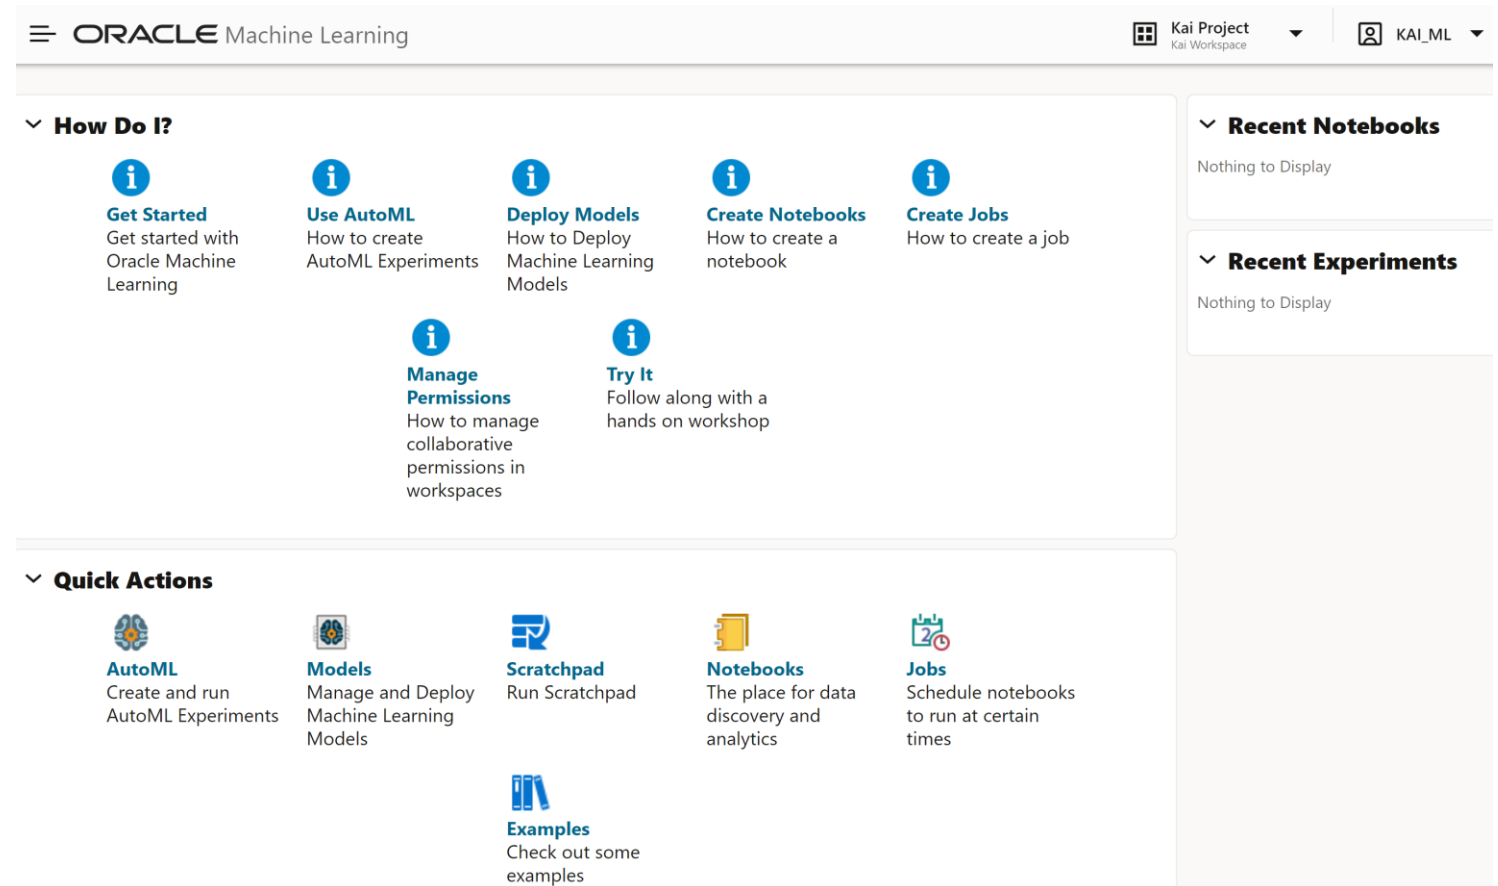
Database name:  
OML1

Sign in with your Oracle Machine Learning Database User credentials

USERNAME  
kai\_ml

PASSWORD  
.....

Sign In



ORACLE Machine Learning

Kai Project  
Kai Workspace

KAI\_ML

### How Do I?

- Get Started**  
Get started with Oracle Machine Learning
- Use AutoML**  
How to create AutoML Experiments
- Deploy Models**  
How to Deploy Machine Learning Models
- Create Notebooks**  
How to create a notebook
- Create Jobs**  
How to create a job
- Manage Permissions**  
How to manage collaborative permissions in workspaces
- Try It**  
Follow along with a hands on workshop

### Quick Actions

- AutoML**  
Create and run AutoML Experiments
- Models**  
Manage and Deploy Machine Learning Models
- Scratchpad**  
Run Scratchpad
- Notebooks**  
The place for data discovery and analytics
- Jobs**  
Schedule notebooks to run at certain times
- Examples**  
Check out some examples

### Recent Notebooks

Nothing to Display

### Recent Experiments

Nothing to Display

# Running Oracle Machine Learning with Autonomous Database

## Machine Learning Example Templates: OML4SQL

The screenshot displays the Oracle Machine Learning interface. At the top, the navigation bar includes the Oracle logo and 'Machine Learning' text on the left, and 'Kai Project' and 'KAI\_ML' on the right. Below the navigation bar, the main content area is titled 'Example Templates'. A search bar on the right contains the text 'oml4sql'. A '+ Create Notebook' button is located on the left. The main area features a grid of 12 notebook templates, each with a title, description, author, date added, tags, and interaction icons (likes, views, shares).

Template Title	Description	Author	Date Added	Tags	Likes	Views	Shares
<b>OML4SQL Anomaly Detection MSET</b>	This notebook illustrates the use of Multivariate Spac...	Oracle	1/6/22, 2:42 PM	'21c' 'SQL' 'Anomaly Detection' 'PREDICTION_D...	0	0	0
<b>OML4SQL Anomaly Detection SVM</b>	This notebook builds an anomaly detection (1-Class-...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'SQL' 'Anomaly Detection' 'Join' 'Sam...	0	0	0
<b>OML4SQL Association Rules Apriori</b>	This notebook builds an associations rules model (A ...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'SQL' 'Association Rules' 'A Priori' 'M...	0	0	0
<b>OML4SQL Attribute Importance MDL</b>	This notebook computes attribute importance (Mini...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'SQL' 'Attribute Importance' 'SH.SUP...	0	0	0
<b>OML4SQL Classification DT</b>	This notebook builds and applies a Decision Tree Cla...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'SQL' 'Classification' 'Split' 'SH.SUPPL...	0	0	0
<b>OML4SQL Classification GLM</b>	This notebook builds and applies a Generalized Line...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'SQL' 'Classification' 'Split' 'SH.SUPPL...	0	0	0
<b>OML4SQL Classification NB</b>	This notebook builds and applies a Naïve Bayes Clas...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'SQL' 'Classification' 'Split' 'SH.SUPPL...	0	0	0
<b>OML4SQL Classification NN</b>	This notebook builds and applies a Neural Network ...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'SQL' 'Classification' 'Split' 'SH.SUPPL...	0	0	0
<b>OML4SQL Classification RF</b>	This notebook builds and applies a Random Forest C...	Oracle					
<b>OML4SQL Classification SVM</b>	This notebook builds and applies a Support Vector ...	Oracle					
<b>OML4SQL Classification XGBoost</b>	This notebook builds and applies an XGBoost Classifi...	Oracle					
<b>OML4SQL Clustering EM</b>	This notebook builds and applies clustering models (...)	Oracle					

# Running Oracle Machine Learning with Autonomous Database

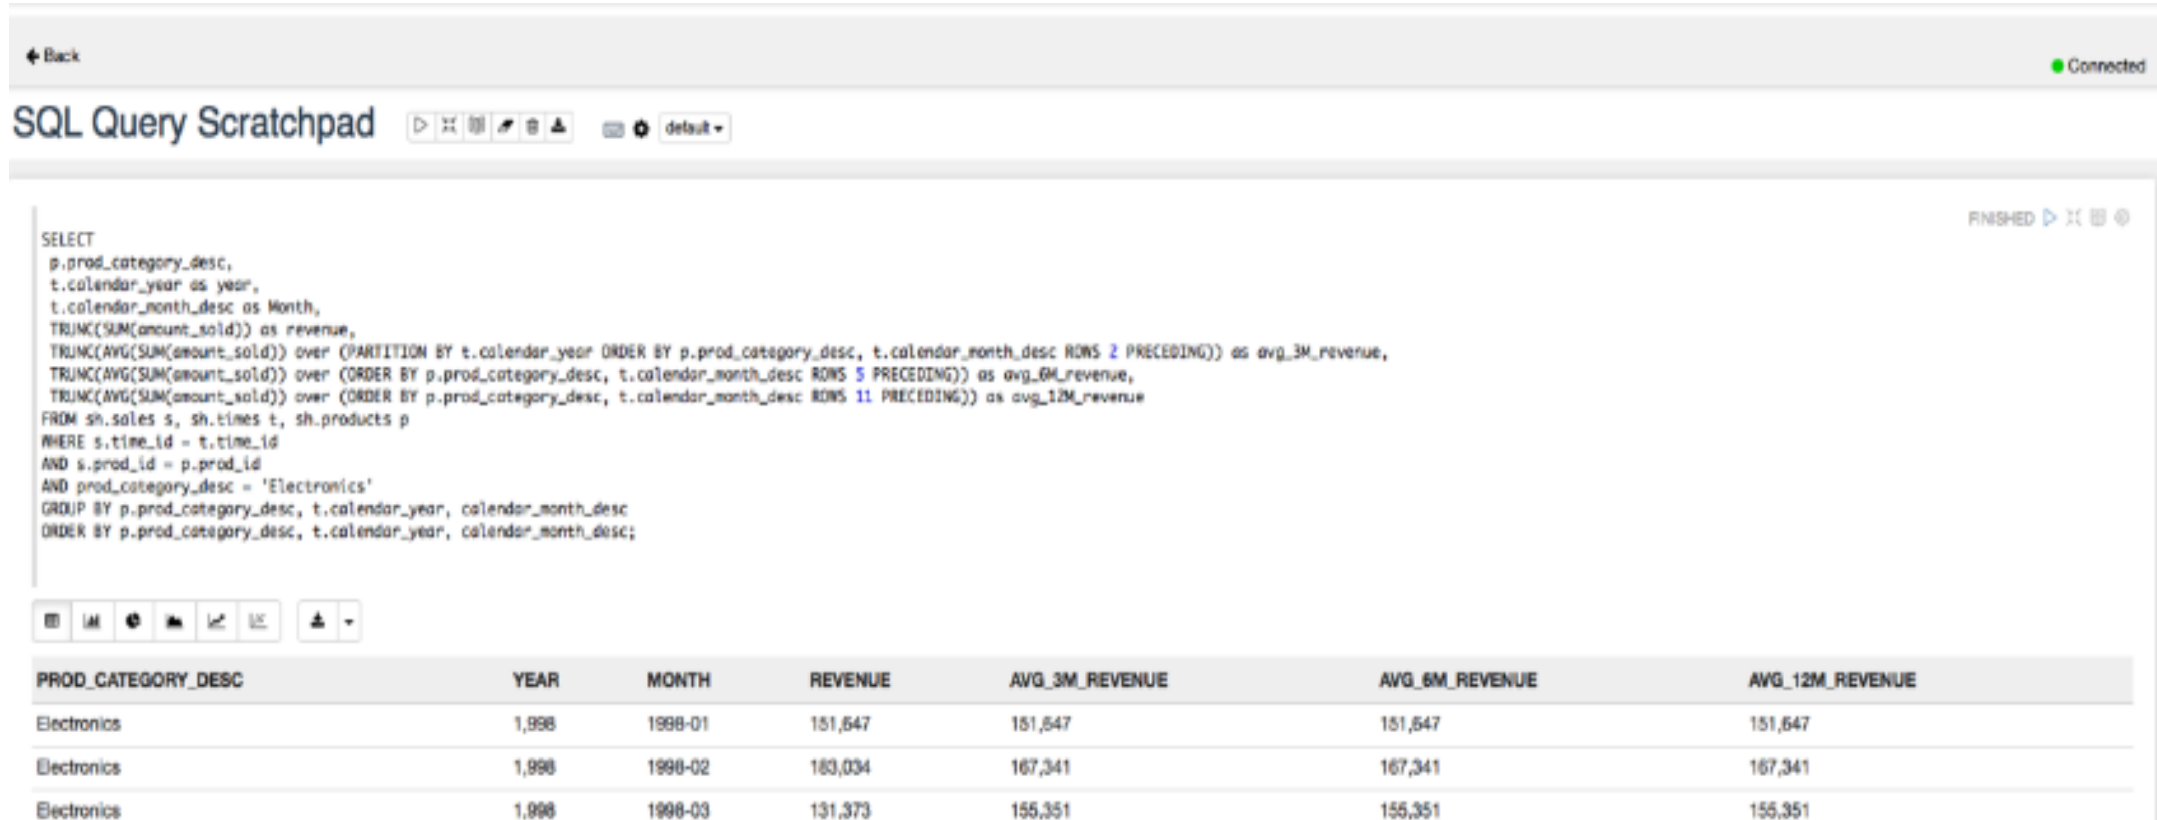
## Machine Learning Example Templates: OML4Py

The screenshot displays the Oracle Machine Learning interface. At the top left, it says "ORACLE Machine Learning". On the top right, there are dropdown menus for "Kai Project" and "KAI\_ML". Below the header, there is a section titled "Example Templates" with a "+ Create Notebook" button. A search bar on the right contains the text "oml4py". The main area is a grid of 12 notebook templates, each with a title, description, author, date added, tags, and interaction icons (likes, views, shares).

Template Title	Description	Author	Date Added	Tags	Likes	Views	Shares
<b>OML4Py -0- Tour</b>	This notebook highlights a wide range of OML4Py fe...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'Python' 'AutoML' 'Summary Statistic...	0	0	0
<b>OML4Py -1- Introduction</b>	This notebook highlights OML4Py core features	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'Python' 'Model Build' 'Model Apply' ...	0	0	0
<b>OML4Py -2- Data Selection and Mani...</b>	This notebook highlights the OML4Py Transparency ...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'Python' 'Correlation' 'Iris' 'Data Selec...	0	0	0
<b>OML4Py -3- Datastore and Script Re...</b>	This notebook highlights features of OML4Py datast...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'Python' 'Embedded Python' 'Iris' 'Da...	0	0	0
<b>OML4Py -4- Embedded Python Exec...</b>	This notebook highlights features of OML4Py Embed...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'Python' 'Embedded Python' 'REST In...	0	0	0
<b>OML4Py -5- AutoML</b>	This notebook highlights the AutoML features of O...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'Python' 'AutoML' 'Random Forest' 'C...	0	0	0
<b>OML4Py Anomaly Detection SVM</b>	This notebook builds an anomaly detection (1-Class-...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'Python' 'Anomaly Detection' 'Join' 'S...	0	0	0
<b>OML4Py Association Rules Apriori</b>	This notebook builds an associations rules model (A ...	Oracle	1/6/22, 2:42 PM	'19c' '21c' 'Python' 'Association Rules' 'A Priori' ...	0	0	0
<b>OML4Py Attribute Importance MDL</b>	This notebooks builds an attribute importance mode...	Oracle					
<b>OML4Py Classification DT</b>	This notebook builds and applies a Decision Tree Cla...	Oracle					
<b>OML4Py Classification GLM</b>	This notebook builds and applies a Generalized Line...	Oracle					
<b>OML4Py Classification NB</b>	This notebook builds and applies a Naive Bayes Clas...	Oracle					

# Running Oracle Machine Learning with Autonomous Database

Use Oracle Machine Learning: Login to Oracle ML, taken into OML home page ;Run SQL Statement in SQL Query Scratchpad



The screenshot displays the Oracle SQL Query Scratchpad interface. At the top left, there is a 'Back' button. At the top right, a green dot indicates 'Connected'. The main title is 'SQL Query Scratchpad' with a 'default' dropdown menu. Below the title, there are several icons for query execution and management. The SQL query is as follows:

```
SELECT
p.prod_category_desc,
t.calendar_year as year,
t.calendar_month_desc as Month,
TRUNC(SUM(amount_sold)) as revenue,
TRUNC(AVG(SUM(amount_sold)) over (PARTITION BY t.calendar_year ORDER BY p.prod_category_desc, t.calendar_month_desc ROWS 2 PRECEDING)) as avg_3M_revenue,
TRUNC(AVG(SUM(amount_sold)) over (ORDER BY p.prod_category_desc, t.calendar_month_desc ROWS 5 PRECEDING)) as avg_6M_revenue,
TRUNC(AVG(SUM(amount_sold)) over (ORDER BY p.prod_category_desc, t.calendar_month_desc ROWS 11 PRECEDING)) as avg_12M_revenue
FROM sh.sales s, sh.times t, sh.products p
WHERE s.time_id = t.time_id
AND s.prod_id = p.prod_id
AND prod_category_desc = 'Electronics'
GROUP BY p.prod_category_desc, t.calendar_year, calendar_month_desc
ORDER BY p.prod_category_desc, t.calendar_year, calendar_month_desc;
```

The query results are displayed in a table with the following columns: PROD\_CATEGORY\_DESC, YEAR, MONTH, REVENUE, AVG\_3M\_REVENUE, AVG\_6M\_REVENUE, and AVG\_12M\_REVENUE. The results show three rows of data for the 'Electronics' category for the years 1998, 1999, and 2000.

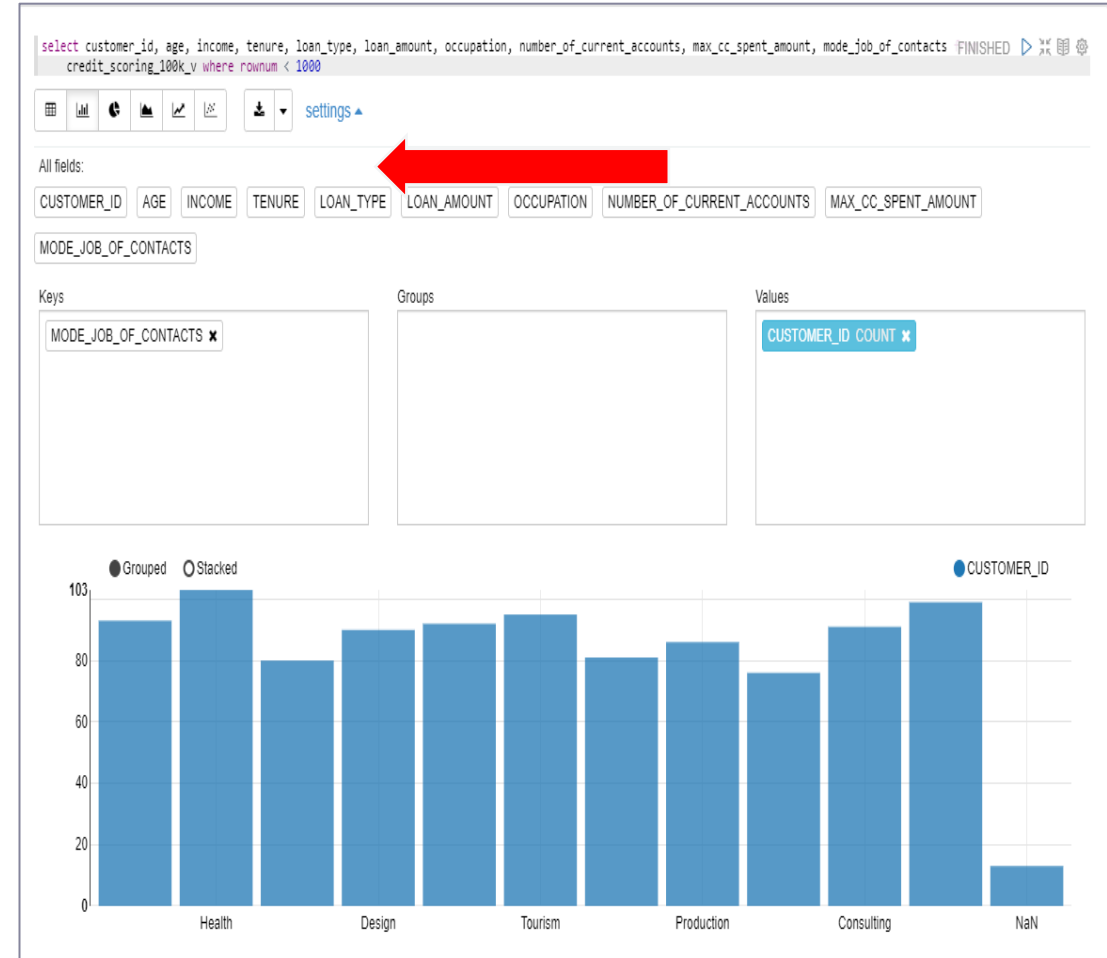
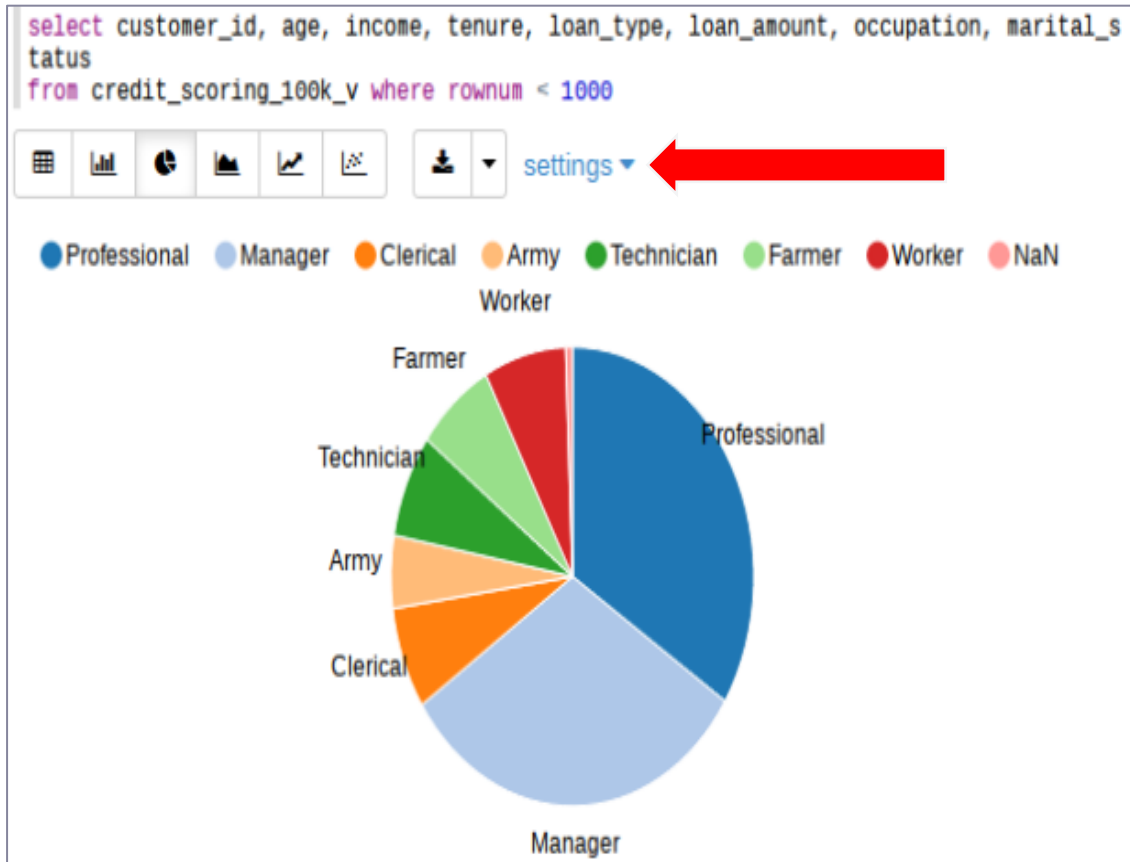
PROD_CATEGORY_DESC	YEAR	MONTH	REVENUE	AVG_3M_REVENUE	AVG_6M_REVENUE	AVG_12M_REVENUE
Electronics	1,998	1998-01	151,647	151,647	151,647	151,647
Electronics	1,998	1998-02	163,034	167,341	167,341	167,341
Electronics	1,998	1998-03	131,373	155,351	155,351	155,351

# Oracle Machine Learning SQL Notebook

The report menu bar lets you change the result to a graph and/or export result

Changing the display

Specify the statistical functions





# Oracle Machine Learning for Python(OML4Py)

OML4Py is a Python module that enables Python users to manipulate data in database tables and views using Python syntax.

Available in two environments:

- In Oracle Autonomous Database, we can use the Python interpreter in Oracle machine Learning notebook.
- In an on-premises database, we can install OML4Py client connecting to OML4Py in database.

need to install Python, the required Python libraries, an OML4Py client, the OML4Py server components in the db

Learn OML4Py with Oracle Autonomous Database through examples

## OML4Py -1- Introducti...



FINISHED

## Oracle Machine Learning for Python (OML4Py)

**Oracle Machine Learning for Python (OML4Py)** makes the open source Python scripting language and environment ready for the enterprise and big data. Designed for problems involving both large and small data volumes, OML4Py integrates Python with Oracle Autonomous Database, allowing users to run Python commands and scripts for statistical, machine learning, and visualization analyses on database tables and views using Python syntax. Many familiar Python functions are overloaded that translate Python behavior into SQL for running in-database, as well as new automatic machine learning capabilities.

In this notebook, we highlight a few of the OML4Py features:

- Load OML library
- Create database tables
- Use the transparency layer
- Rank attributes for predictive value using the in-database attribute importance algorithm
- Build predictive models
- Score data using these models

Copyright (c) 2021 Oracle Corporation

[The Universal Permissive License \(UPL\), Version 1.0](#)

# Oracle Machine Learning for Python(OML4Py)

Import libraries supporting OML4Py

```
%python

import warnings
warnings.filterwarnings('ignore')

import pandas as pd
import oml
from oml import automl

pd.set_option('display.max_rows', 500)
pd.set_option('display.max_columns', 500)
pd.set_option('display.width', 1000)
```

Took 0 secs. Last updated by ML\_USER at December 18 2021, 10:34:01 PM.

Creating a new table with ***oml.create*** based on data from sklearn datasets.

```
FINISHED ▶
%python
from sklearn import datasets
import pandas as pd
pd.set_option('display.width',1000)

breastcancer = datasets.load_breast_cancer()
bc_data = breastcancer.data.astype(float)
X = pd.DataFrame(bc_data, columns = breastcancer.feature_names)
y = pd.DataFrame(breastcancer.target, columns = ['TARGET'])

try:
    oml.drop(table='BREASTCANCER')
except:
    pass

BREASTCANCER = oml.create(pd.concat([X, y], axis=1), table =
    'BREASTCANCER') # returns proxy object

print("Shape:",BREASTCANCER.shape)
print("Description:\n", breastcancer.DESCR)
```

# An Example: Building Machine Learning Model with OML

# Building Machine Learning Model with OML

ORACLE Machine Learning

ADWC\_WS Project [ADWC\_WS Workspace]

ADWC\_WS

Connected

## Credit Score Predictions

FINISHED

### Targeting Likely Good Credit Customers using Oracle Machine Learning's (OML) Classification Models

Heather has spent most of her time over the past couple of years extracting and preparing data for analysis. The large volumes of data need extracting and processing mean she spends most of her time waiting for jobs to finish and very little of her time analyzing the data. Demands from marketing are forcing a new approach whereby the data remains in the data warehouse and is processed there. The alternative cloud solution is more complex, and has no direct out of the box processes to analyze the data in place. She started taking a look at Oracle, and found the simple SQL commands in ADWC are familiar, and execute extremely fast, leveraging all the performance features of the platform. Further once she is done can can apply the learning models to incoming data on the fly, and allow end user analysts to immediately see mining results. This drastically reduces the cycle of data preparation, analysis, and publishing. It also means there is no change to analysis/reporting Data Visualization toolset that users are familiar with.

Scroll down this notebook and learn how to use OML to create predictive perspectives on data in ADWC, WITHOUT moving it. We will process a small 100k data set, but could use a 100M or billion row data set without worrying about processing time.


This is an extract of Alphaoffice customer information. We will first get acclimated to Apache Zeppelin, the open source interface for interactive collaboration in a team environment.

### The Business Problem:

#### Increase Sales by Targeting our Best Customers; Good Credit Customers!

Heather has a hunch that weakening sales may be due to the company selling to non-optimal customers; customers who perhaps have poor credit and fail to make their payments for their purchases. Heather has over 100 variables to consider so wants to first explore her data using simple charts and graphs, but then move onto using Oracle Machine Learning's powerful algorithms to automatically sift through her data to find patterns, new insights and to make predictions that target her best customers—those who have good credit.

Took 1 sec. Last updated by ADWC\_WS at August 17 2018, 1:14:54 PM.



FINISHED

Took 1 sec. Last updated by ADWC\_WS at July 12 2018, 3:40:22 PM. (outdated)

# Building Machine Learning Model with OML

## OML Example: predicate customer credit score:

Based on existing customer information with known credit\_sore in customer table

- Customer\_id
- Age
- Income
- Tenure
- Loan\_type
- Loan\_amount
- Occupation
- Marital\_status
- Wealth
- Credit\_limit
- ....
- 
- Credit\_score\_bin

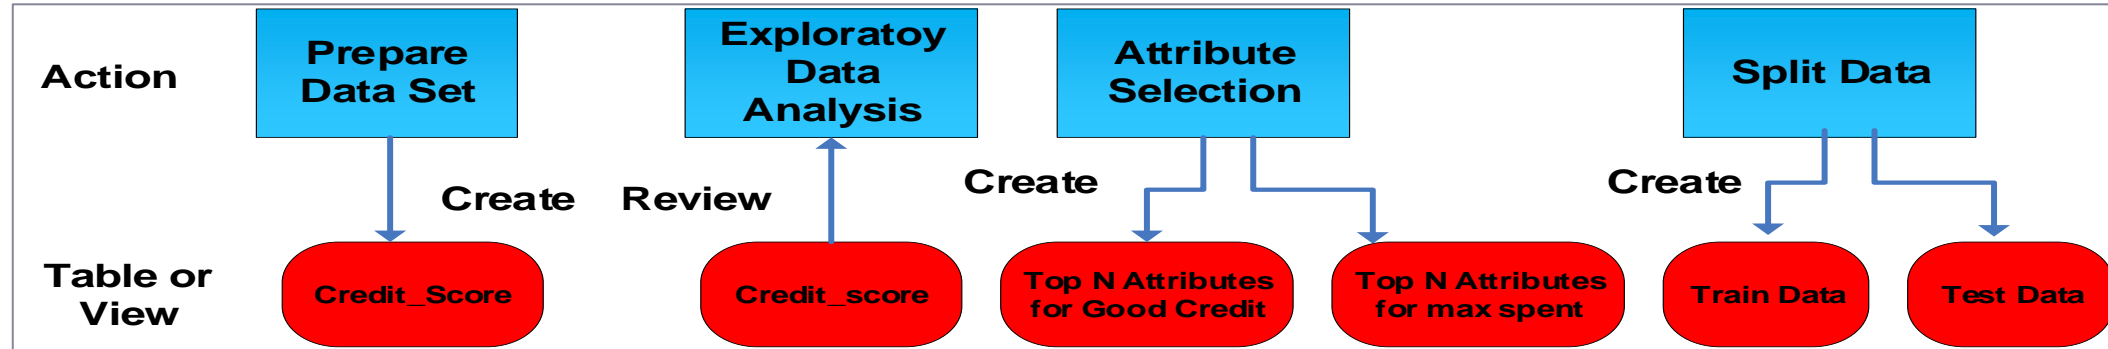
To create a machine learning model that can predict the potential credit scores for new customers whose credit scores are unknown

You can try this by yourself by following this link:

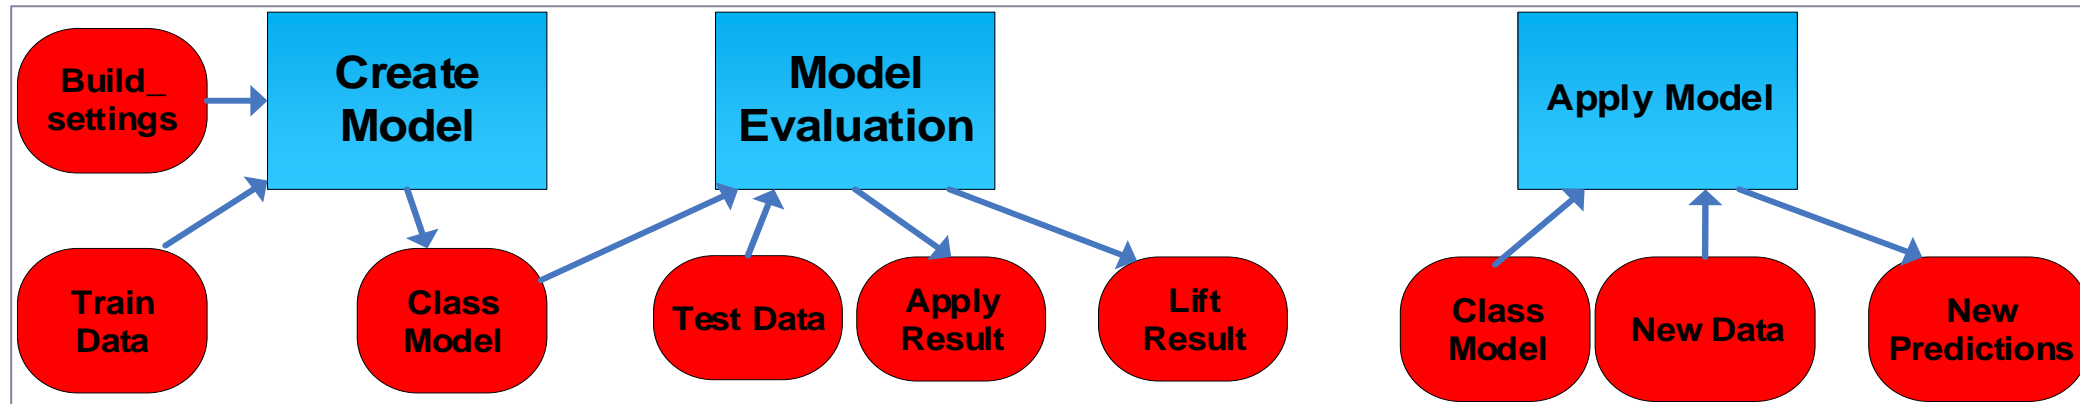
<https://github.com/oracle/learning-library/blob/master/workshops/adwc4dev/L300.md>

# Building Machine Learning Model with OML

## Data Gathering and Preparation



## ML Model Building, Evaluation and Application



# Data Gathering and Preparation with OML

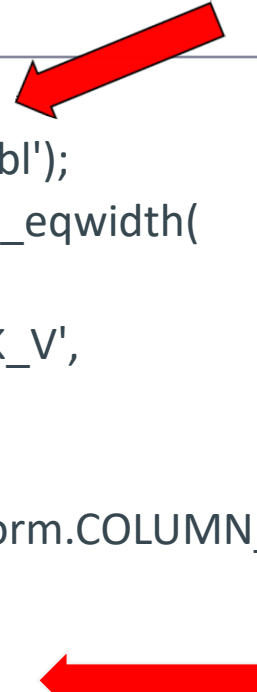
Prepare Data set:

```
create view credit_scoring_100k as select * from credit_score
```

Exploratory Data Analysis: bin the variable


PL/SQL code:

```
dbms_data_mining_transform.create_bin_num(  
    bin_table_name => 'bin_num_tbl');  
dbms_data_mining_transform.insert_autobin_num_eqwidth(  
    bin_table_name => 'bin_num_tbl',  
    data_table_name => 'CREDIT_SCORING_100K_V',  
    bin_num        => 5,  
    max_bin_num    => 10,  
    exclude_list   => dbms_data_mining_transform.COLUMN_LIST('CUSTOMER_ID'));  
dbms_data_mining_transform.xform_bin_num(  
    bin_table_name => 'bin_num_tbl',  
    data_table_name => 'CREDIT_SCORING_100K_V',  
    xform_view_name => 'mining_data_bin_view');
```

Two red arrows are present in the image. One arrow points from the top right towards the function name 'create\_bin\_num' in the first line of the PL/SQL code. The second arrow points from the right towards the function name 'xform\_bin\_num' in the last line of the PL/SQL code.

# Data Gathering and Preparation with OML

## Exploratory Data Analysis: bin the variable

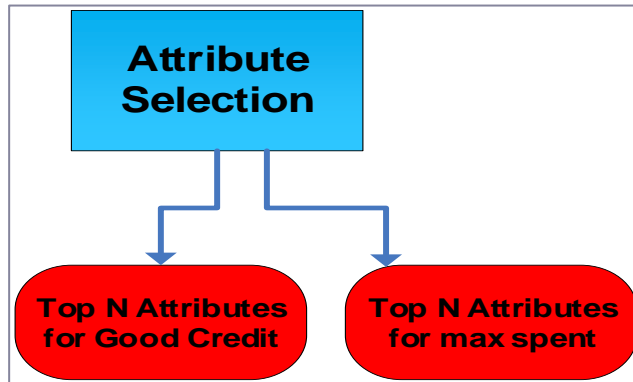




# Data Gathering and Preparation with OML

## Attribute Selection aka Feature Engineering/Feature Selection

Create attribute importance Machine Learning model for Good Credit



```
%script
DECLARE
v_sql varchar2(100);

BEGIN
BEGIN EXECUTE IMMEDIATE 'DROP TABLE ai_explain_output_credit_score_bin';
EXCEPTION WHEN OTHERS THEN NULL;
END;

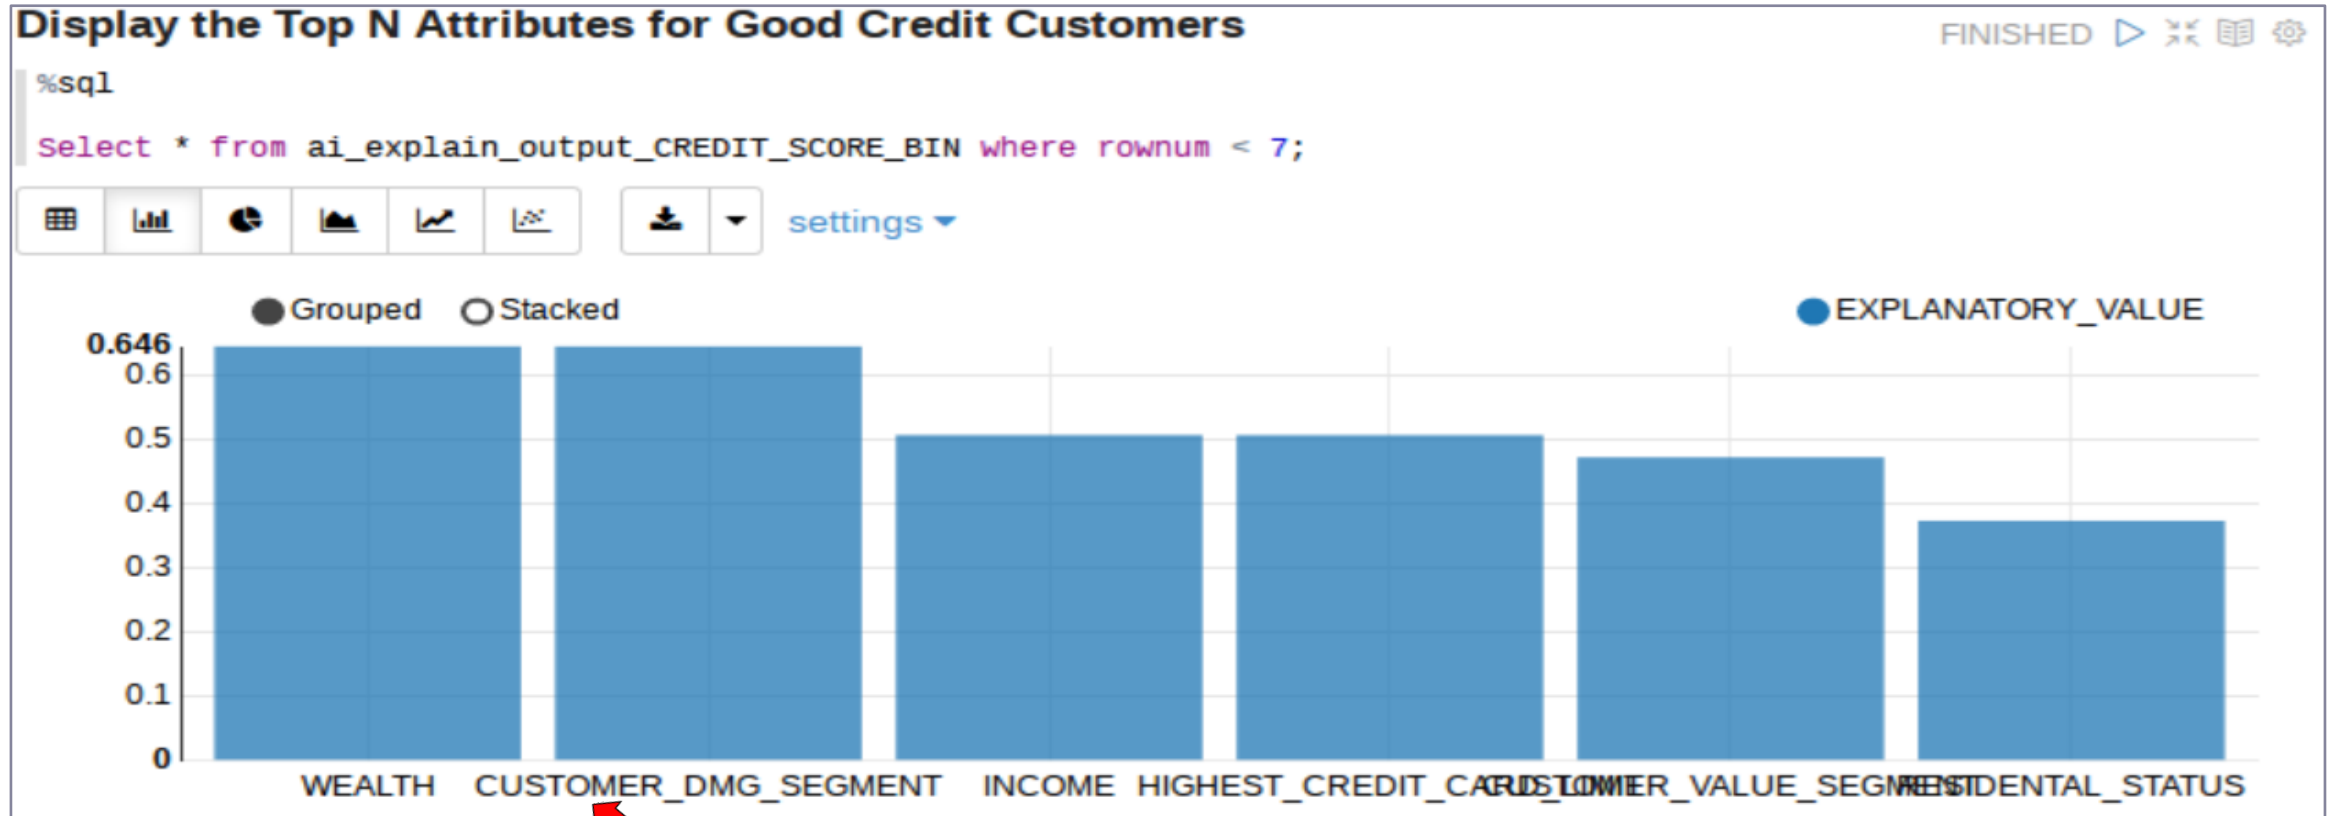
BEGIN
DBMS_PREDICTIVE_ANALYTICS.EXPLAIN(
  data_table_name      => 'CREDIT_SCORING_100K_V',
  explain_column_name => 'CREDIT_SCORE_BIN',
  result_table_name   => 'AI_EXPLAIN_OUTPUT_CREDIT_SCORE_BIN');
END;
End;
```

PL/SQL procedure successfully completed.

# Data Gathering and Preparation with OML

Attribute Selection aka Feature Engineering/Feature Selection

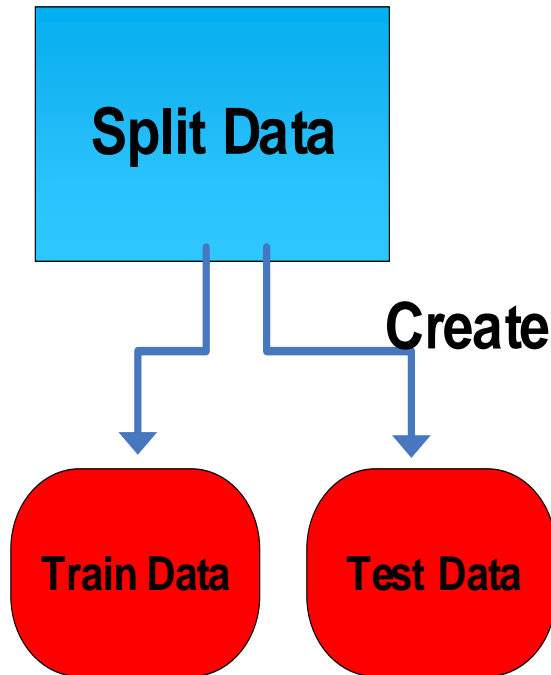
Create attribute importance Machine Learning model for Good Credit



Top attributes for good credit customer

# Data Gathering and Preparation with OML

- . Split Data into Train data (60% data): N1\_TRAIN\_DATA  
Test Data (40% data) : N1\_TEST\_DATA




Create or replace view N1\_TRAIN\_DATA  
as SELECT \* from CREDIT\_scoring\_100k\_V SAMPLE (60) SEED (1)

Create or replace view N2\_TEST\_DATA  
as SELECT \* from CREDIT\_scoring\_100k MINUS  
SELECT \* FROM N1\_TRAIN\_DATA

# Model Building, Evaluation and Application with OML

## Create Predictive Model to Target Good Credit Customers

- Add parameter settings to n1\_build\_settings table



```
/* Create a Build Setting (DT) for Model Build */
```

```
EXECUTE IMMEDIATE 'CREATE TABLE n1_build_settings (setting_name VARCHAR2(30),setting_value VARCHAR2(4000))';
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings (setting_name, setting_value) VALUES (''ALGO_NAME'', ''ALGO_DECISION_TREE'')';
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings (setting_name, setting_value) VALUES (''PREP_AUTO'', ''ON'')';
```

```
DBMS_OUTPUT.PUT_LINE ('Created model build settings table: n1_build_settings ');
```

```
/*
```

```
-- Populate and Adjust Model Setting (DT) for Model Build
```

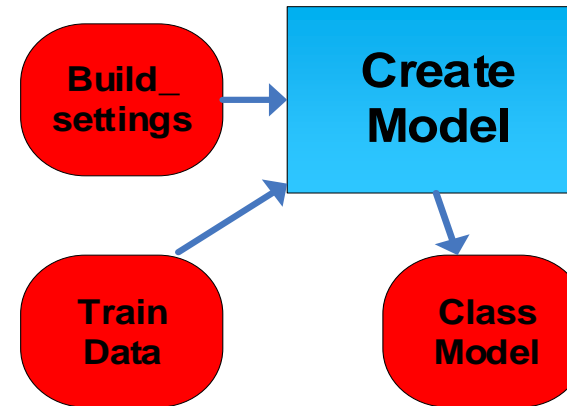
```
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings VALUES (''TREE_TERM_MAX_DEPTH'', 7)';
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings VALUES (''TREE_TERM_MINREC_SPLIT'', 20)';
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings VALUES (''TREE_TERM_MINPCT_SPLIT'', .1)';
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings VALUES (''TREE_TERM_MINREC_NODE'', 10)';
EXECUTE IMMEDIATE 'INSERT INTO n1_build_settings VALUES (''TREE_TERM_MINPCT_NODE'', 0.05)';
```

```
*/
```

# Model Building, Evaluation and Application with OML

## Create Predictive Model to Target Good Credit Customers

- Build a Classification model



```
BEGIN
DBMS_DATA_MINING.CREATE_MODEL(
  model_name           =>'N1_CLASS_MODEL',
  mining_function      =>'dbms_data_mining.classification',
  data_table_name      =>'N1_TRAIN_DATA',
  case_id_column_name  =>'CUSTOMER_ID',
  target_column_name   =>'CREDIT_SCORE_BIN',
  setting_table_name   =>'N1_BUILD_SETTINGS')
END;
```

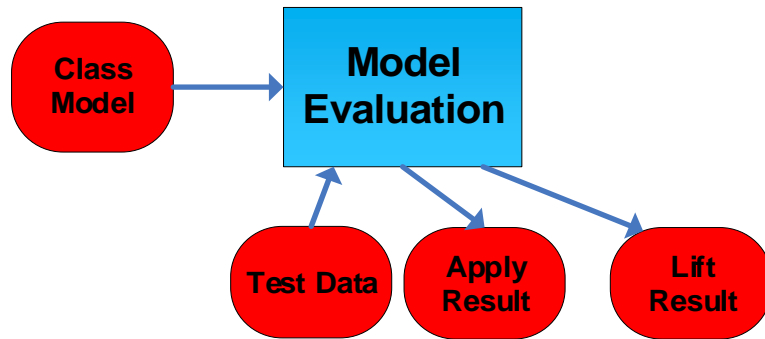
# Model Building, Evaluation and Application with OML

## Model Evaluation

Test the Model by generating an apply result

Execute DBMS\_DATA\_MINING.APPLY Procedure:

```
DBMS_DATA_MINING.APPLY('N1_CLASS_MODEL','N1_TEST_DATA','CUSTOMER_ID','N1_APPLY_RESULT')
```



Test Data: N1\_TEST\_DATA  
Apply Result: N1\_APPLY\_RESULT  
Lift Result: N1\_LIFT\_TABLE

- Create a lift result : Execute DBMS\_DATA\_MINING.COMPUTE\_LIFT Procedure:

```
DBMS_DATA_MINING.COMPUTE_LIFT ('N1_APPLY_RESULT', 'N1_TEST_DATA', 'CUSTOMER_ID', 'CREDIT_SCORE_BIN',  
'N1_LEFT_TABLE', 'GOOD CREDIT', 'PREDICATION', 'PROBABILITY', 100)
```



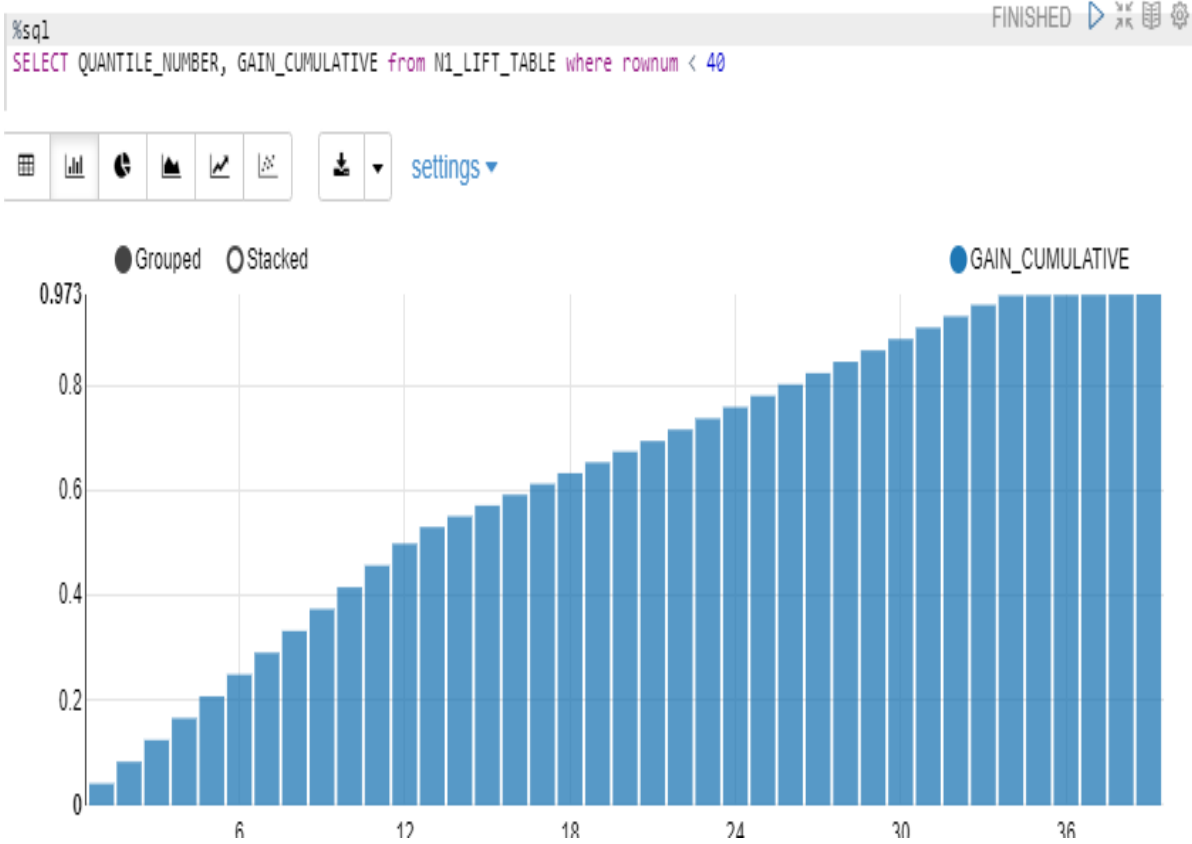
# ML Model Building, Evaluation and Application with OML

## Model Evaluation

- Review Lift result :View Model's Cumulative Gains Chart and decide if its a good model

### Gains and Lift Charts


- Gain or lift is a measure of the effectiveness of a classification model calculated as the ratio between the results obtained with and without the model
- Gain and lift charts are visual aids for evaluating performance of classification models
- Both charts consist of a lift curve and a baseline
- However, in contrast to the confusion matrix that evaluates models on the whole population gain or lift chart evaluates model performance in a portion of the population
- The greater the area between the lift curve and the baseline, the better the model!



# Model Building, Evaluation and Application with OML

## Apply Model to New data:

Apply the Oracle Machine Learning Model to New Customers to Show Customers Most Likely to Have Good Credit  
(for new customers in credit\_scoring\_new\_cust\_v)



**%sql**


```
select a.customer_id , a.prob_Credit_Score_Bin , b.age, b.income, b.tenure, b.loan_type, b.loan_amount,
b.occupation, b.education_level, b.marital_status
from (select * from (select Customer_id, round(prob_Credit_Score_Bin *100,2)
                prob_Credit_Score_Bin
                from (select Customer_ID, prediction_probability(N1_CLASS_MODEL,
                NULL using *) prob_Credit_Score_Bin from
                credit_scoring_new_cust_v))) a,
        credit_scoring_100k_v b
where a.customer_id = b.customer_id order by a.prob_Credit_Score_Bin desc
```



# Model Building, Evaluation and Application with OML

Apply Model to New data:

Put the predication results into  
credit\_score\_new\_predictions table



```
%  
create table credit_score_new_predictions as  
select a.customer_id  
      , a.prob_good_credit  
      , b.age, b.income, b.tenure, b.loan_type, b.loan_amount, b.occupation, b.education_level, b.marital_status  
from (select * from (select Customer_id, round(prob_good_credit *100,2) prob_good_credit from (select Customer_ID,  
      prediction_probability(N1_CLASS_MODEL, 'Good Credit' using *) prob_good_credit from credit_scoring_new_cust_v))) a  
      , credit_scoring_100k_v b  
where a.customer_id = b.customer_id
```

FINISHED ▶ ✖ 📖 ⚙️

Updated 9920 row(s).

# Model Building, Evaluation and Application with OML

Apply Model to New data:

Apply a ML Model to a Single Record in a Transactional Application:

```
select prediction_probability(N1_CLASS_MODEL, 'Good Credit' USING 'Rich' as WEALTH, 2000 as income, 'Silver' as customer_value_segment) Prediction_Probability from dual;
```

```
%  
select prediction_probability(N1_CLASS_MODEL, 'Good Credit'  
    USING 'Rich' as WEALTH, 2000 as income, 'Silver' as customer_value_segment) Prediction_Probability  
from dual;
```



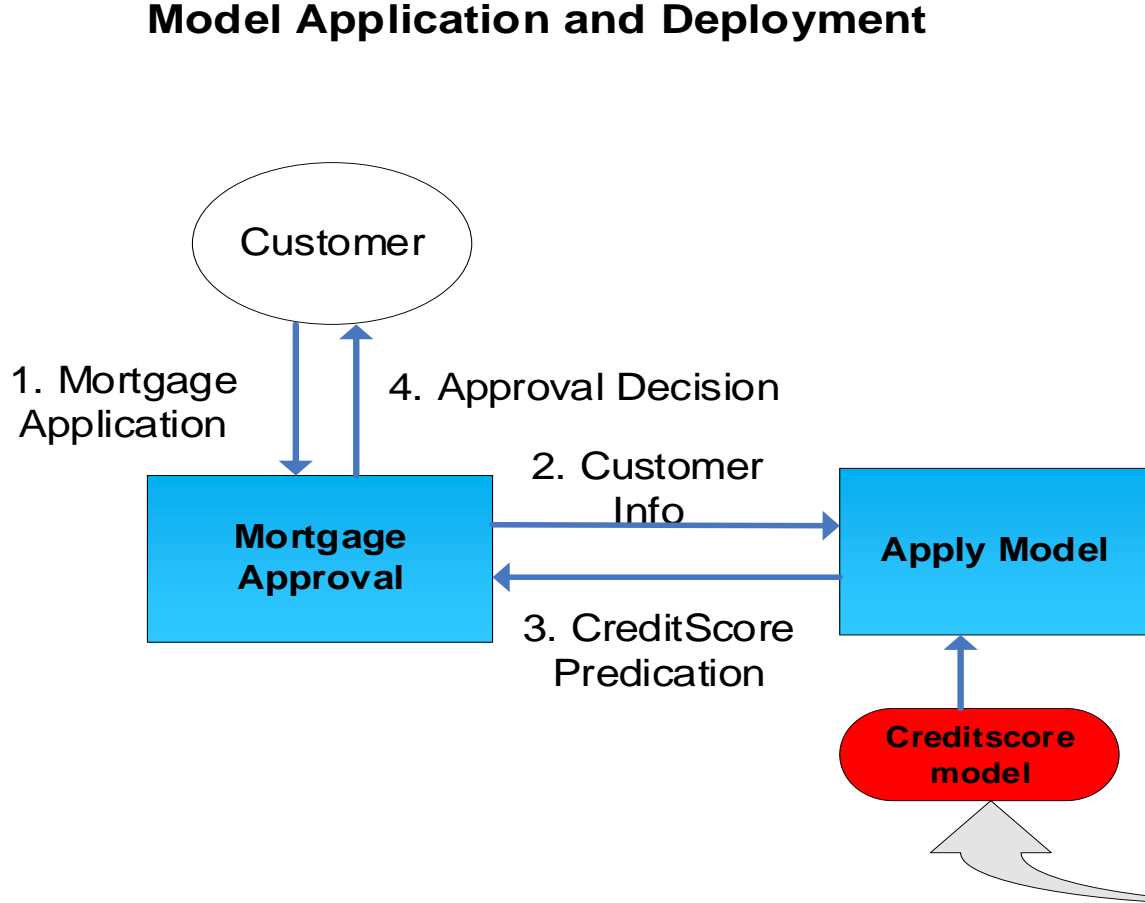
**PREDICTION\_PROBABILITY**

0.5098856966326846

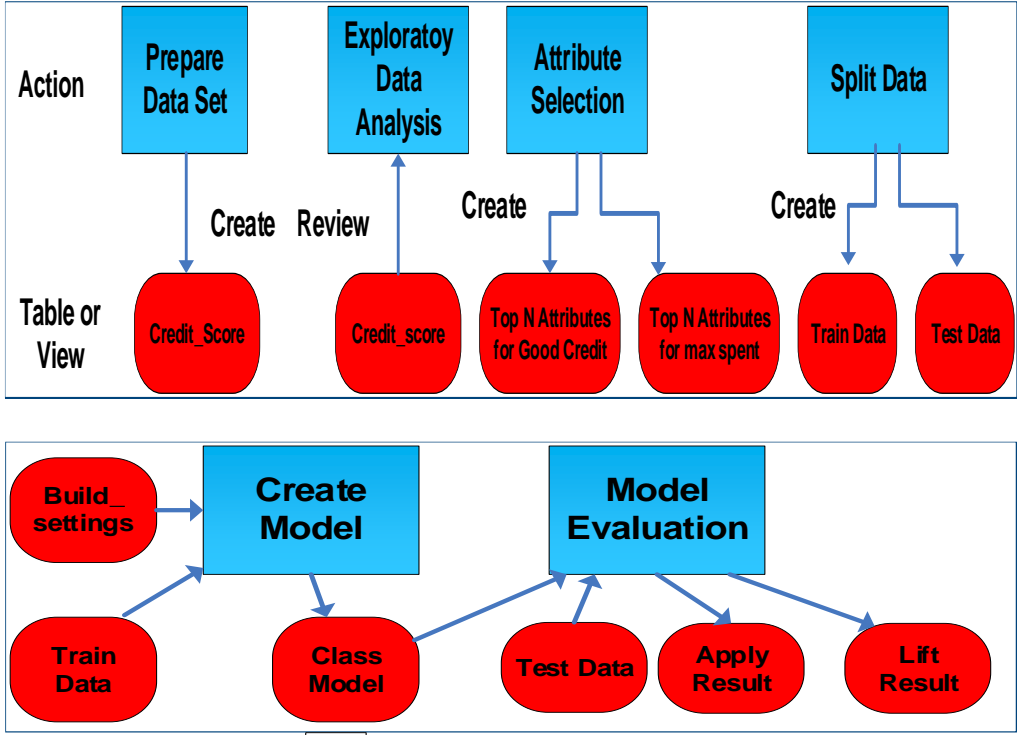
---

# Enabling Business Applications with machine Learning

An example: Using Machine Learning for a mortgage approval application




### Data preparation and Model Development



# Summary and QA

- Machine Learning Enabling Business Applications
- Machine Learning in Oracle Database
  - Move Algorithm, Not Data
  - Oracle Machine Learning for SQL (OML4SQL)
  - Oracle Machine Learning for Python (OML4Py)
- Running Oracle Machine Learning with Oracle Autonomous Database
  - Machine Learning Algorithms in Oracle Autonomous Database
  - SQL Query Scratchpad
  - Oracle Machine Learning SQL Notebook
  - OML4SQL and OML4Py in Oracle Autonomous Database
- An Example: Building a Machine Learning Model with OML in Autonomous Database
  - Data Gathering and Preparation
  - ML Model Building, Evaluation and Application

# Related Publication



## About the Authors



**Heli Helskyaho** is the CEO of Miracle Finland Oy. She holds a master's degree in computer science from the University of Helsinki and specializes in databases. She is currently working on her doctoral studies, researching and teaching at the University of Helsinki. Her research areas cover big data, multi-model databases, schema discovery, and methods and tools for utilizing semi-structured data for decision making.

Heli has been working in IT since 1990. She has held several positions, but every role has included databases and database designing. She believes that understanding your data makes using the data much easier. She is an Oracle ACE Director, an Oracle Groundbreaker Ambassador, and a

frequent speaker at many conferences. She is the author of several books and has been listed as one of the top 100 influencers in the IT sector in Finland for each year from 2015 to 2020.

Heli can be reached at [www.linkedin.com/in/helihelskyaho/](https://www.linkedin.com/in/helihelskyaho/), <https://helifromfinland.blog>, and <https://twitter.com/HeliFromFinland>.



**Jean Yu** is a Senior Staff MLOps Engineer at Habana Labs, an Intel company. Prior to that, she was a Senior Data Engineer on the IBM Hybrid Cloud Management Data Science Team. Her current interests include deep learning, model productization, and distributed training of massive transformer-based language models. She holds a master's degree in computer science from the University of Texas at San Antonio. She has more than 25 years of experience in developing, deploying, and managing software applications,

### ABOUT THE AUTHORS

as well as in leading development teams. Her recent awards include an Outstanding Technical Achievement Award for significant innovation in Cloud Brokerage Cost and Asset Management products in 2019 as well as an Outstanding Technical Achievement Award for Innovation in the Delivery of Remote Maintenance Upgrade for Tivoli Storage Manager in 2011.

Jean is a master inventor with 14 patents granted. She has been a voting member of the IBM Invention Review Board from 2006 to 2020. She has been a speaker at conferences such as North Central Oracle Apps User Group Training Day 2019 and Collaborate 2020.

Jean can be reached at [www.linkedin.com/in/jean-yu/](https://www.linkedin.com/in/jean-yu/).



**Kai Yu** is a Distinguished Engineer in the Dell Technical Leadership Community. He is responsible for providing technical leadership to Dell Oracle Solutions Engineering. He has more than 27 years of experience in architecting and implementing various IT solutions, specializing in Oracle databases, IT infrastructure, the cloud, business analytics, and machine learning.

Kai has been a frequent speaker at various IT/Oracle conferences, with over 200 presentations in more than 20 countries. He has authored 36 articles in technical journals such as *IEEE Transactions on Big Data* and co-authored the

book *Expert Oracle RAC 12c* (Apress, 2013). He has been an Oracle ACE Director since 2010. He has served on the IOUG/Quest Conference committee, as the IOUG RAC SIG president, and as the IOUG CLOUD SIG co-founder and vice president. He received the 2011 OAUG Innovator of Year award and the 2012 Oracle Excellence Award for Technologist of the Year: Cloud Architect by *Oracle Magazine*. He holds master's degrees in computer science and engineering from the Huazhong University of Science and Technology and the University of Wyoming.

Kai can be reached at <https://kyuoracleblog.wordpress.com>, [www.linkedin.com/in/kaiyu1](https://www.linkedin.com/in/kaiyu1), and [https://twitter.com/ky\\_austin1](https://twitter.com/ky_austin1).

**Thank You!**

kai.yu@dell.com

<http://kyuoracleblog.wordpress.com/>