



ANALYTICS AND DATA

TechCasts

Build Practical AI Prototypes with Oracle Analytics (Even with a Small Team)

Dennis Gray, Data Migration Specialist, Definian
Adam Hume, Senior Consultant, Definian

Past & Future TechCasts:



February 5th

Our Favorite Features of OAC:
November 2025 & January 2026 Releases

Presented by Tim Vlamis, Branden Pavol, Wayne Van Sluys, Gautam Pisharam & Taiwo Ajayi



March 5th

Fraud Stops with Oracle Database 26ai
Vector Search + OML

Presented by Abi Giles-Haigh



April 2nd

Build Practical AI Prototypes with Oracle
Analytics (Even with a Small Team)

Presented by Dennis Gray & Adam Hume

TechCast Archive

2026	2025	2024	2023	2022	2021
Date	Title	Presenter(s)		Replay	Download(s)
Mar 5	Fraud Stops with Oracle Database 26ai Vector Search + OML	Dr. Abigail Giles-Haigh		Video	Slides
Feb 5	Our Favorite Features of OAC: November 2025 & January 2026 Releases	Branden Pavol, Gautam Pisharam, & Carter Beaton		Video	Slides
Jan 8	The Oracle AI Microservices Sandbox for RAG Rapid Prototyping	Melli Annamalai & Phani Chilakapati		Video	Slides

Submit a topic to share at <https://andouc.org/techcasts/>

Let's Connect



Website

<http://andouc.org/>



Chat with the Experts

<https://bit.ly/Join-ANDOUC-Slack>



Watch Previous TechCasts

<https://bit.ly/3qmGgHN>




<https://www.linkedin.com/company/analytics-and-data-oracle-user-community>



Spatial + Graph SIG

bit.ly/Spatial-Graph-LinkedIn



@analyticsndataouc.
bsky.social

Feedback Panel

Oracle PM's:



Emcees:




- Provide feedback directly to Oracle PM's
- Find out what's coming next
- Tell them what you'd like to see next in future releases

Thursday - 1:30p



Analytics and Data

SUMMIT 2026




Only 1 week to go!
April 14-16, 2026



Scan this for full information:





Build Practical AI Prototypes with Oracle Analytics (Even with a Small Team)

A disciplined approach to launching AI
Presented by Dennis Gray and Adam Hume

Spring 2026

What We Mean by “Prototype”

(Not an Experiment, not a Demo)

A prototype should:

- Be explainable
- Be reproducible
- Fit into existing architecture
- Be expandable into production

AI Prototypes Don't Fail Because of Models

They Fail Because of Poor Foundations

Small teams rarely struggle with:

- Access to algorithms
- Access to models
- Access to tutorials

They struggle with:

- Architecture
- Data discipline
- Repeatability
- Governance clarity

01

Rules for Prototypes

The following rules are not about slowing innovation down. They are about making sure your prototype can survive beyond a demo and evolve into something durable.

Rule #1: Structure before Speed


Govern Naming and Structure Early

Govern:

- Folder hierarchy
- Dataset naming conventions
- Model naming conventions
- Data flow naming
- Output labeling

Why?

- Enhances discoverability
- Makes workflow self-documenting
- Enables task delegation
- Prevents “mystery artifacts”
- Reduces rework



Rule #2: Safe Data before Sensitive Data

Prototype with Public Data

Use:

- Public datasets (data.gov, USDA, etc.)
- Synthetic enterprise-like datasets
- Smaller volume subsets

Why?

- Prevents accidental data leakage
- Avoids premature compliance overhead
- Reduces cost exposure
- Encourages experimentation



<https://catalog.data.gov/dataset/>



<https://www.kaggle.com/datasets>



<https://www.chatgpt.com>

Rule #3: Architecture before Algorithms


Pipes Before Predictions

Before model selection:

- Where will data live?
- How will it refresh?
- Who owns it?
- How will results surface?
- How will results be consumed?

The most sophisticated model in the world is useless if:

- It cannot be refreshed
- It cannot be explained
- It cannot be embedded into dashboards



02

Foundation

Every practical AI solution starts with a secure foundation. In OCI, that foundation is networking, identity, and policy control.

Orchestrating OCI Services for Practical AI

A secure, scalable foundation

Infrastructure Foundation

- VCN (Virtual Cloud Network) *
- IAM (Identity & Access Management)
- Policies & Dynamic Groups
- Object Storage *

Data Layer

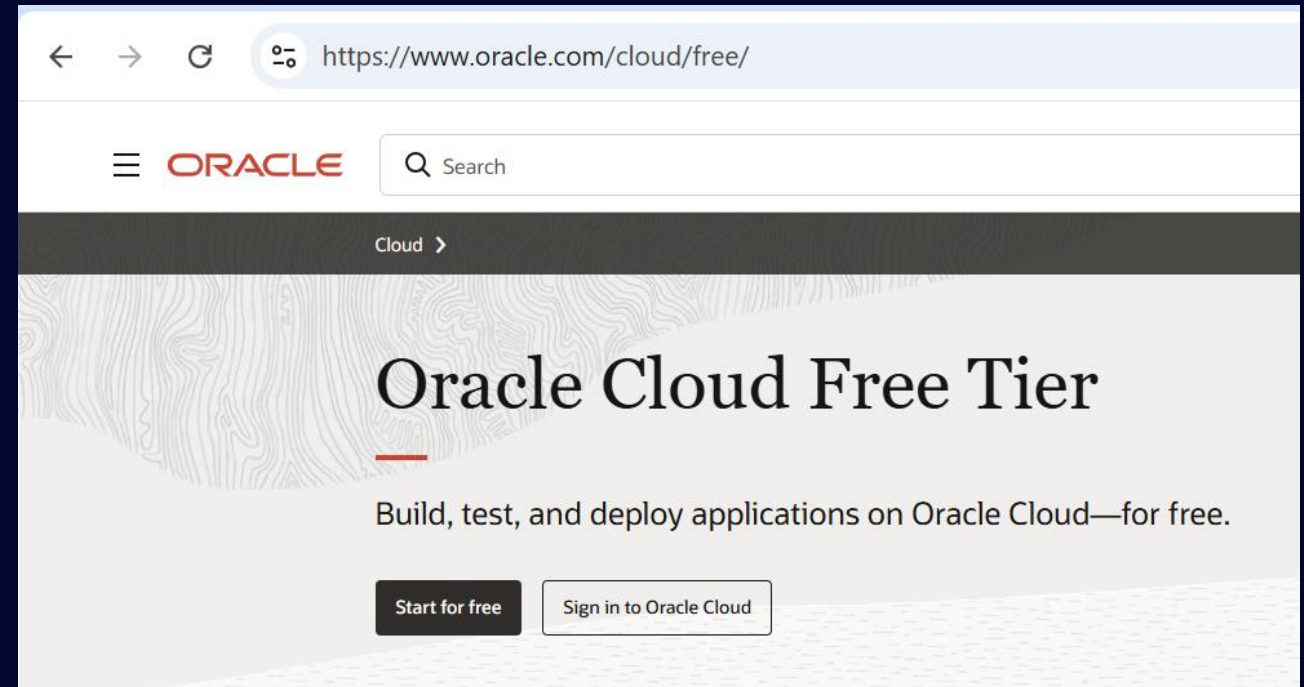
- **Oracle Autonomous Database 26ai** *
- Stores structured data, vectors, embeddings

Modeling Layer

- **Oracle Cloud Infrastructure Data Science**
- Builds, trains, and deploys ML models

Insight Layer

- **Oracle Analytics Cloud**
- Delivers dashboards, AI explanations, and business insights



Virtual Cloud Network (VCN)

The Foundation of Secure AI in Oracle Cloud Infrastructure

A Virtual Cloud Network (VCN) is your private, software-defined network inside OCI.

It defines:

- IP address ranges (CIDR blocks)
- Subnets (public & private)
- Route tables (traffic direction)
- Gateways (Internet, NAT, Service)
- Security rules (firewall controls)

Why My AI Prototype Needed It

- Before I launch my Data Science instance:
- It had to live inside a network boundary
- OCI needed to assign private IPs
- Access to Object Storage and databases had to be controlled
- Security policies had to be enforced

Even with “default settings,” I’m operating inside enterprise-grade infrastructure!

The screenshot shows the OCI console interface for Virtual Cloud Networks. The left sidebar contains a navigation menu with 'Networking' selected. The main content area is titled 'Virtual Cloud Networks' and includes a search bar, an 'Applied filters' section showing 'Compartment definianflorida (root)', and a 'Create VCN' button. Below these elements is a table with columns for 'Name', 'State', and 'IPv4 CIDR Block'. The table is currently empty, displaying the message 'No items to display' and a link to 'Create new items or search again using different filters or search terms.'





Dynamic Groups

Secure Resource-to-Resource Access in Oracle Cloud Infrastructure

A Dynamic Group allows OCI resources (not users) to be granted permissions automatically.

Instead of assigning policies to a person...

We assign policies to cloud resources like:

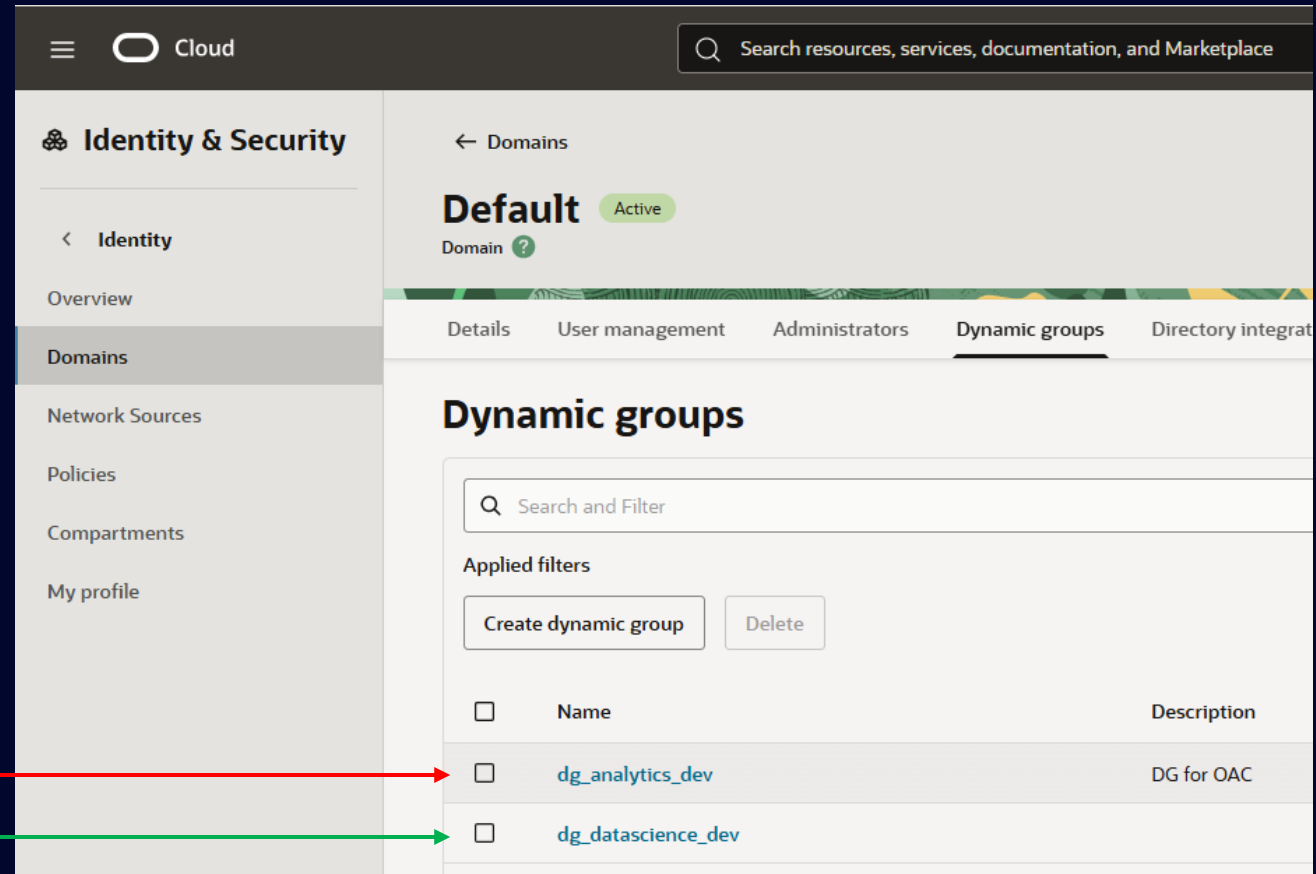
-  Data Science notebook sessions
-  Autonomous Databases
-  Compute instances
-  Analytics instances

Why This Matters for AI Prototypes

When my notebook needs to:

- Read data from Object Storage
- Access a database
- Call other OCI services

It does not use my personal credentials!



The screenshot shows the Oracle Cloud Identity & Security console. The left sidebar is under 'Identity & Security' with 'Domains' selected. The main area shows the 'Default' domain with tabs for 'Details', 'User management', 'Administrators', 'Dynamic groups', and 'Directory integrations'. The 'Dynamic groups' tab is active, displaying a search bar and a table of dynamic groups. Two groups are listed: 'dg_analytics_dev' and 'dg_datascience_dev'. A red arrow points from the 'dg_analytics_dev' group to the red code block below, and a green arrow points from the 'dg_datascience_dev' group to the green code block below.

```
ALL {resource.type = 'analyticsinstance'}
```





```
ALL {resource.type = 'datasciencenotebooksession'}
```

Policies

Granting Access in OCI

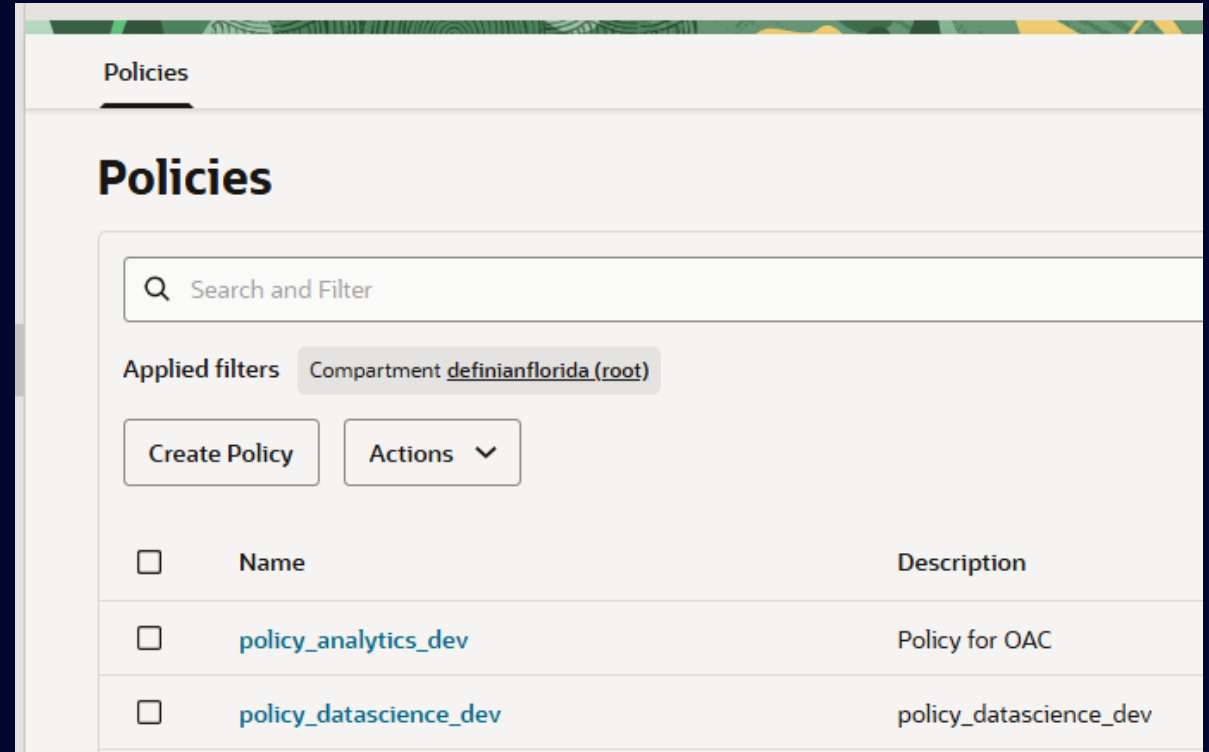
A Policy is a plain-English statement that defines Who can do what to which resources and where.

Policies connect:

-  Users or Groups
-  Dynamic Groups (resources)
-  OCI Services
-  Compartments or Tenancy

Policy Structure:

Allow + subject + to + action + resource-type + in + location



The screenshot shows the OCI Policies console. At the top, there's a 'Policies' header. Below it, a search bar labeled 'Search and Filter' is present. Underneath, 'Applied filters' shows 'Compartment definianflorida.(root)'. There are two buttons: 'Create Policy' and 'Actions' with a dropdown arrow. Below these is a table with columns 'Name' and 'Description'. The table lists two policies: 'policy_analytics_dev' with description 'Policy for OAC' and 'policy_datascience_dev' with description 'policy_datascience_dev'. Each row has a checkbox on the left.

```
Allow dynamic-group dg_analytics_dev to read objectstorage-namespaces in tenancy
Allow dynamic-group dg_analytics_dev to manage object-family in compartment
datascience
```

```
Allow group data_science to use virtual-network-family in tenancy
Allow group data_science to manage virtual-network-family in tenancy
Allow dynamic-group dg_datascience_dev to manage data-science-family in tenancy
Allow dynamic-group dg_datascience_dev to read objectstorage-namespaces in tenancy
Allow dynamic-group dg_datascience_dev to manage object-family in tenancy
```

03

Data Pipeline

Transforming publicly available nutrition data into a structured, analytics-ready dataset inside Oracle 26ai through a simple, repeatable workflow.

From Open Data to Analytics

Building a Nutrition Dataset in Oracle 26ai

- Starting point: USDA Nutrient Database
 - <https://catalog.data.gov/dataset/usda-national-nutrient-database-for-standard-reference-legacy-release-d1570>

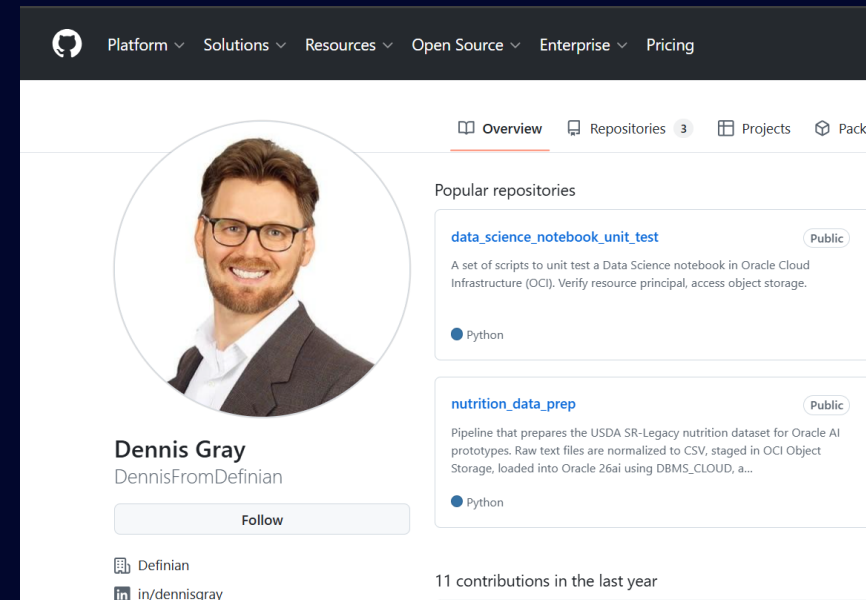
Goal: Prepare analytics-ready data in Oracle 26ai

Workflow overview:

- Download raw data
- Normalize locally
- Upload to OCI Object Storage
- Load into Oracle 26ai
- Model for analytics



https://github.com/DennisFromDefinian/nutrition_data_prep



The screenshot shows the GitHub profile page for Dennis Gray (DennisFromDefinian). The profile includes a circular profile picture, the name "Dennis Gray", and the GitHub handle "DennisFromDefinian". Below the name is a "Follow" button. The "Popular repositories" section lists two repositories: "data_science_notebook_unit_test" and "nutrition_data_prep". The "nutrition_data_prep" repository is highlighted and described as a pipeline that prepares the USDA SR-Legacy nutrition dataset for Oracle AI prototypes. The profile also shows "11 contributions in the last year" and social media links for "Definian" and "in/dennisgray".



Step 1: Downloading the Source Data

Understanding the USDA SR-Legacy Dataset

Dataset characteristics:

- Distributed as multiple relational tables
- ASCII .txt format
- Fields separated by ^
- Text values wrapped in ~

Example tables:

- FOOD_DES (master list of foods and their descriptions)
- FD_GROUP (categories of foods such as dairy, fruits, meats)
- NUT_DATA (nutrient values for each food item)

USDA provided documentation explains:

- Table relationships and keys
- Nutrient identifiers (203 = Protein, 208 = Energy)
- Units of measurement
- Data derivation methods
- Data source references



Step 2: Normalize the Raw Data

Preparing USDA Files for Database Ingestion

Challenge:

The USDA SR-Legacy dataset is not delivered in standard CSV format.

The raw files use:

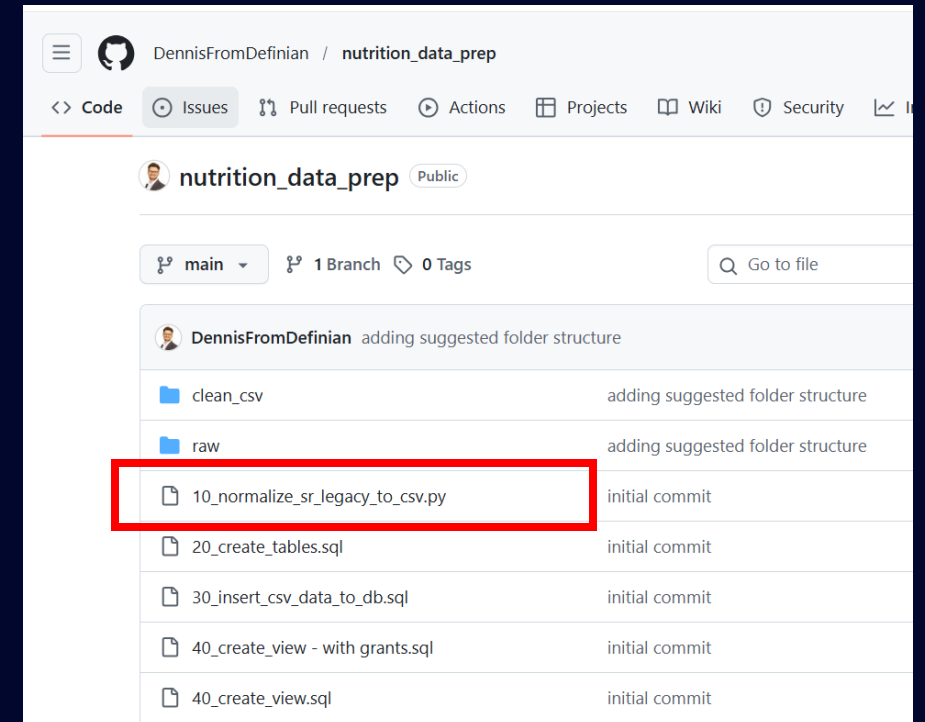
- ^ as a field delimiter
- ~ as a text qualifier
- mixed encoding formats

Transformation Steps

- Convert ^ delimiters → standard CSV format
- Remove ~ text qualifiers
- Normalize encoding to UTF-8
- Preserve original table structure

Python tools used: pandas, csv, pathlib

 Full normalization script available in the project **GitHub repository**



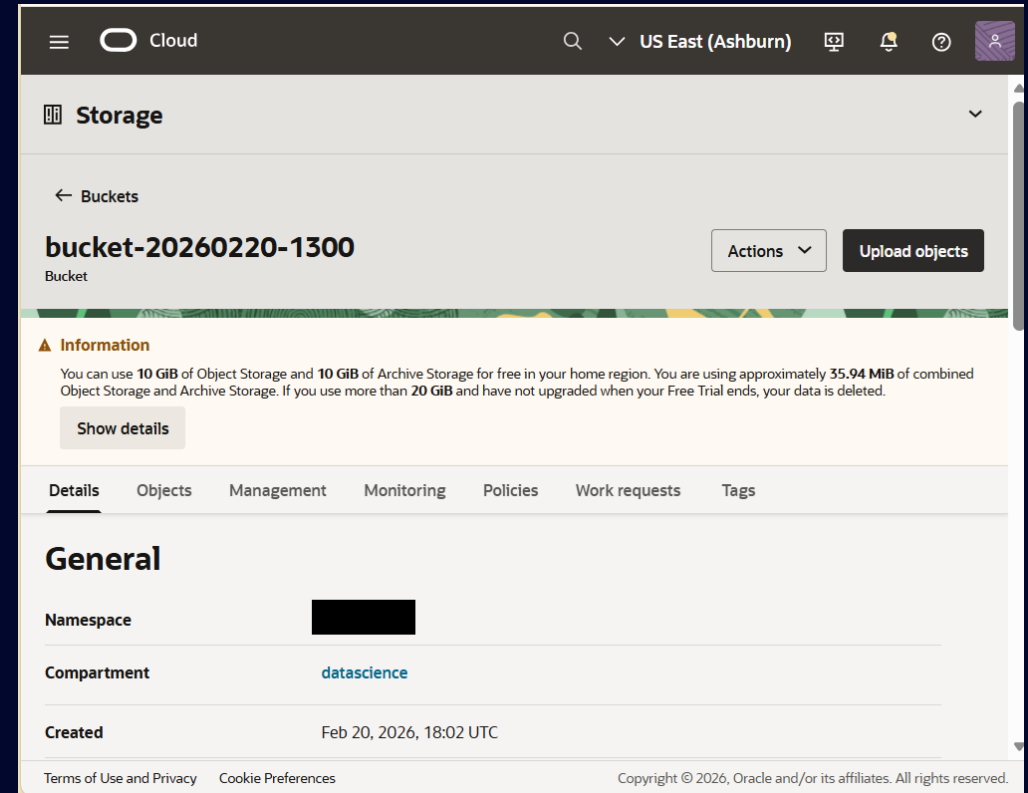
☁ Step 3: Uploading Data to OCI Object Storage

Using OCI Object Storage as a Data Landing Zone

Normalized CSV files uploaded to OCI Object Storage bucket

Benefits:

- Durable staging area
- Easily accessible by Oracle 26ai
- Scales for large datasets





Step 4: Creating Table Structures in Oracle 26ai

Reproducing the USDA Data Model in Oracle 26ai


Input: Normalized CSV files from the USDA SR-Legacy dataset

Design source: USDA SR-Legacy Documentation and User Guide

The documentation provides the schema needed to recreate the tables, including:

- Primary keys and relationships between tables
- Field names and data types
- Field widths and numeric precision
- Nullability rules
- Units of measure and code definitions

Goal: Create Oracle tables that mirror the USDA schema so that the CSV files can be loaded directly without transformation.

```
nutrition_data_prep / 20_create_tables.sql  
  
Code Blame 327 lines (281 loc) · 14.5 KB   
  
59 -- Table 7: Nutrient Definition (NUTR_DEF)  
60 -- Key: Nutr_No  
61 CREATE TABLE NUTR_DEF (  
62     Nutr_No  VARCHAR2(3) NOT NULL,  
63     Units    VARCHAR2(7) NOT NULL,  
64     Tagname  VARCHAR2(20) NULL,  
65     NutrDesc VARCHAR2(60) NOT NULL,  
66     Num_Dec  NUMBER(1,0) NOT NULL,  
67     SR_Order NUMBER(6,0) NOT NULL,  
68     CONSTRAINT PK_NUTR_DEF PRIMARY KEY (Nutr_No)  
69 );  
70
```

Step 5: Load Data from Object Storage

Securely Ingest Data into Oracle 26ai Using OCI Resource Principal

Goal: Load normalized CSV files from OCI Object Storage into Oracle tables

Oracle can load files directly from Object Storage.

This is done using the `DBMS_CLOUD.COPY_DATA` procedure.

Security model

- Access to Object Storage is controlled by OCI IAM policies.
- Earlier we configured Dynamic group for the Autonomous Database and policies allowing the DB to read Object Storage

```
nutrition_data_prep / 30_insert_csv_data_to_db.sql  
  
Code Blame 72 lines (61 loc) · 2.14 KB  
  
28  
29     DBMS_CLOUD.COPY_DATA(  
30         table_name      => p_table,  
31         credential_name => 'OCI$RESOURCE_PRINCIPAL',  
32         file_uri_list  => uri(p_file),  
33         format         => json_object(  
34             'type' value 'csv',  
35             'delimiter' value ',',  
36             'quote' value '',  
37             'skipheaders' value 0,  
38             'blankasnull' value true,  
39             'trimspaces' value 'ltrim',  
40             'ignoremissingcolumns' value true,  
41             'truncatecol' value true,  
42             'conversionerrors' value 'store_null'  
43         )  
44     );
```

Step 6: Modeling the Data for Analytics

Transforming a Normalized Dataset into an Analytics-Friendly Model



The USDA dataset is designed as a normalized relational model. This structure:

- ✓ reduces redundancy
- ✓ preserves data integrity
- ✓ supports maintainability

But it also means that answering simple questions requires multiple joins.

To simplify analysis, we create an analytics view. The resulting dataset provides:

- Food name
- Food category
- Nutritional attributes

```
nutrition_data_prep / 40_create_view.sql   
  
DennisFromDefinian initial commit  
  
Code Blame 58 lines (50 loc) · 3.08 KB   
  
1 CREATE OR REPLACE VIEW VW_FACT_FOOD_NUTRIENTS_100G AS  
2 SELECT  
3     fd.ndb_no,  
4     fd.shrt_desc AS food_item,  
5     fg.fmgrp_desc AS category,  
6  
7     /* Basis for all nutrient values from NUT_DATA:  
8        Nutr_Val is reported per 100g edible portion (SR-Legacy standard). */  
9     'per 100g edible portion' AS basis,  
10
```

Optimized for:

- ✓ Oracle Analytics Cloud
- ✓ Machine learning models
- ✓ clustering and classification
- ✓ AI exploration



Analytics-Ready Nutrition Dataset

A Secure, Enterprise-Ready Dataset for Analytics and Machine Learning

Pipeline summary:

- 1 Download open data
- 2 Normalize locally
- 3 Upload to Object Storage
- 4 Create Oracle tables
- 5 Import data with DBMS_CLOUD
- 6 Model data for analytics

Outcome:

A clean, query-able dataset ready for AI, ML, and analytics.

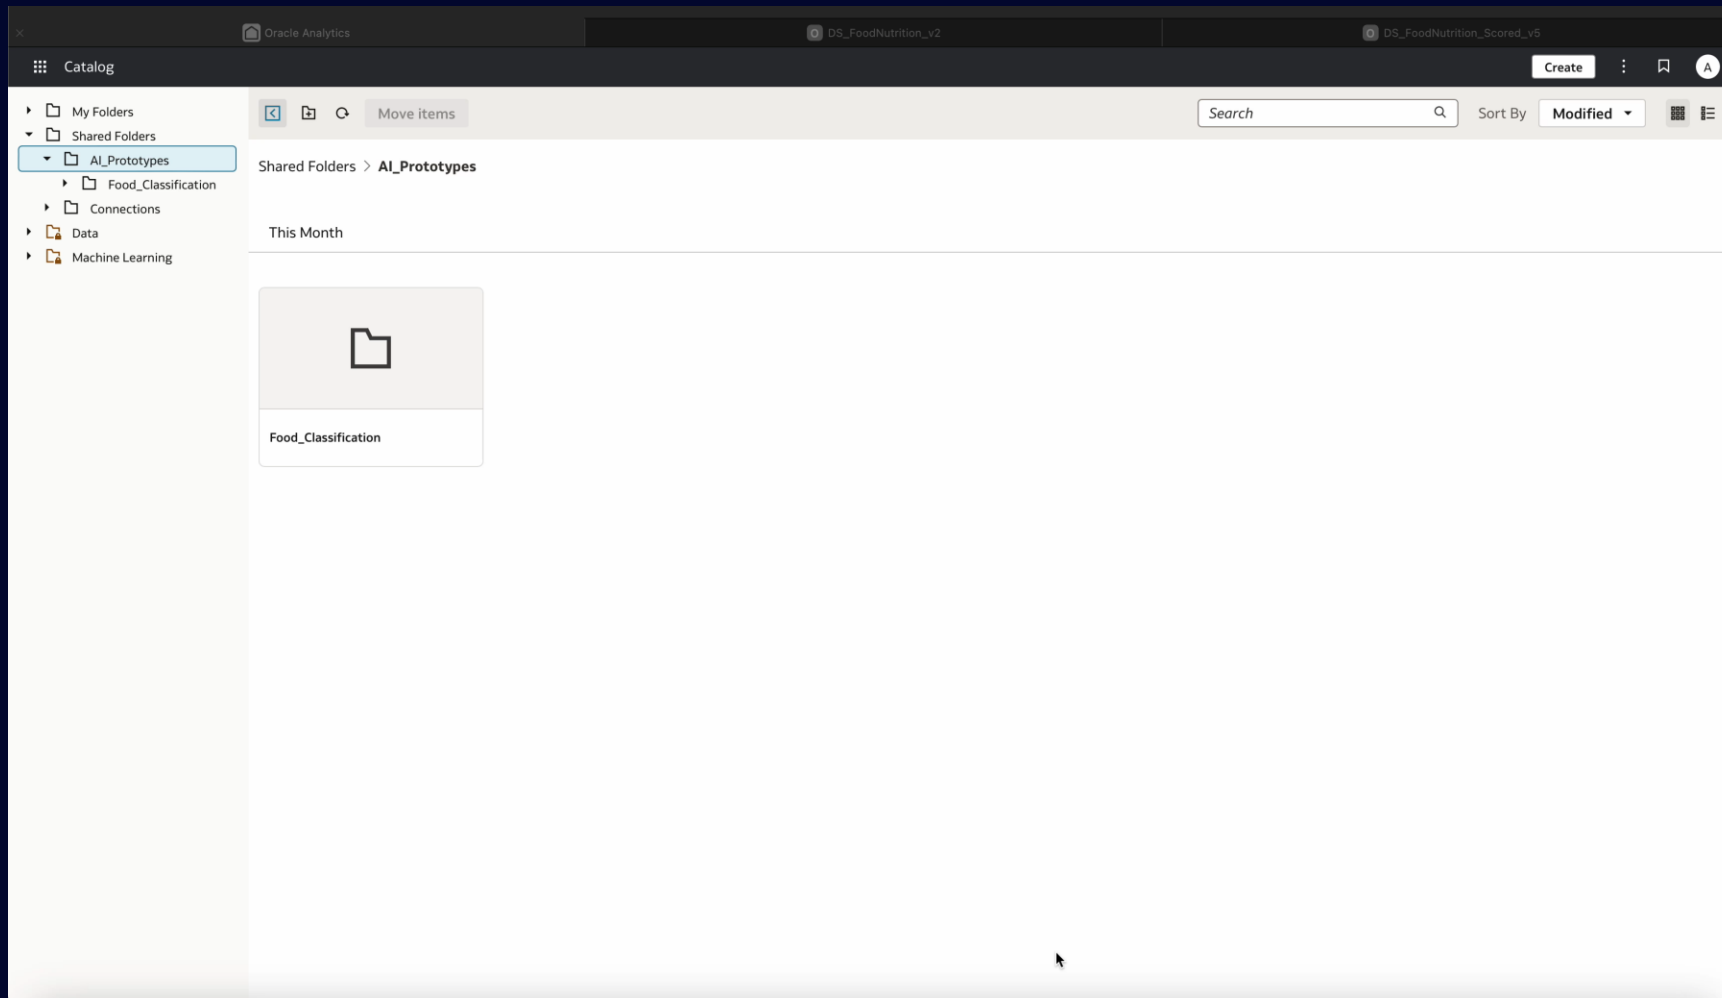
04

Analytics

With the curated dataset in place, we can now use Oracle Analytics Cloud to train machine learning models that identify patterns in the nutritional profiles of foods..

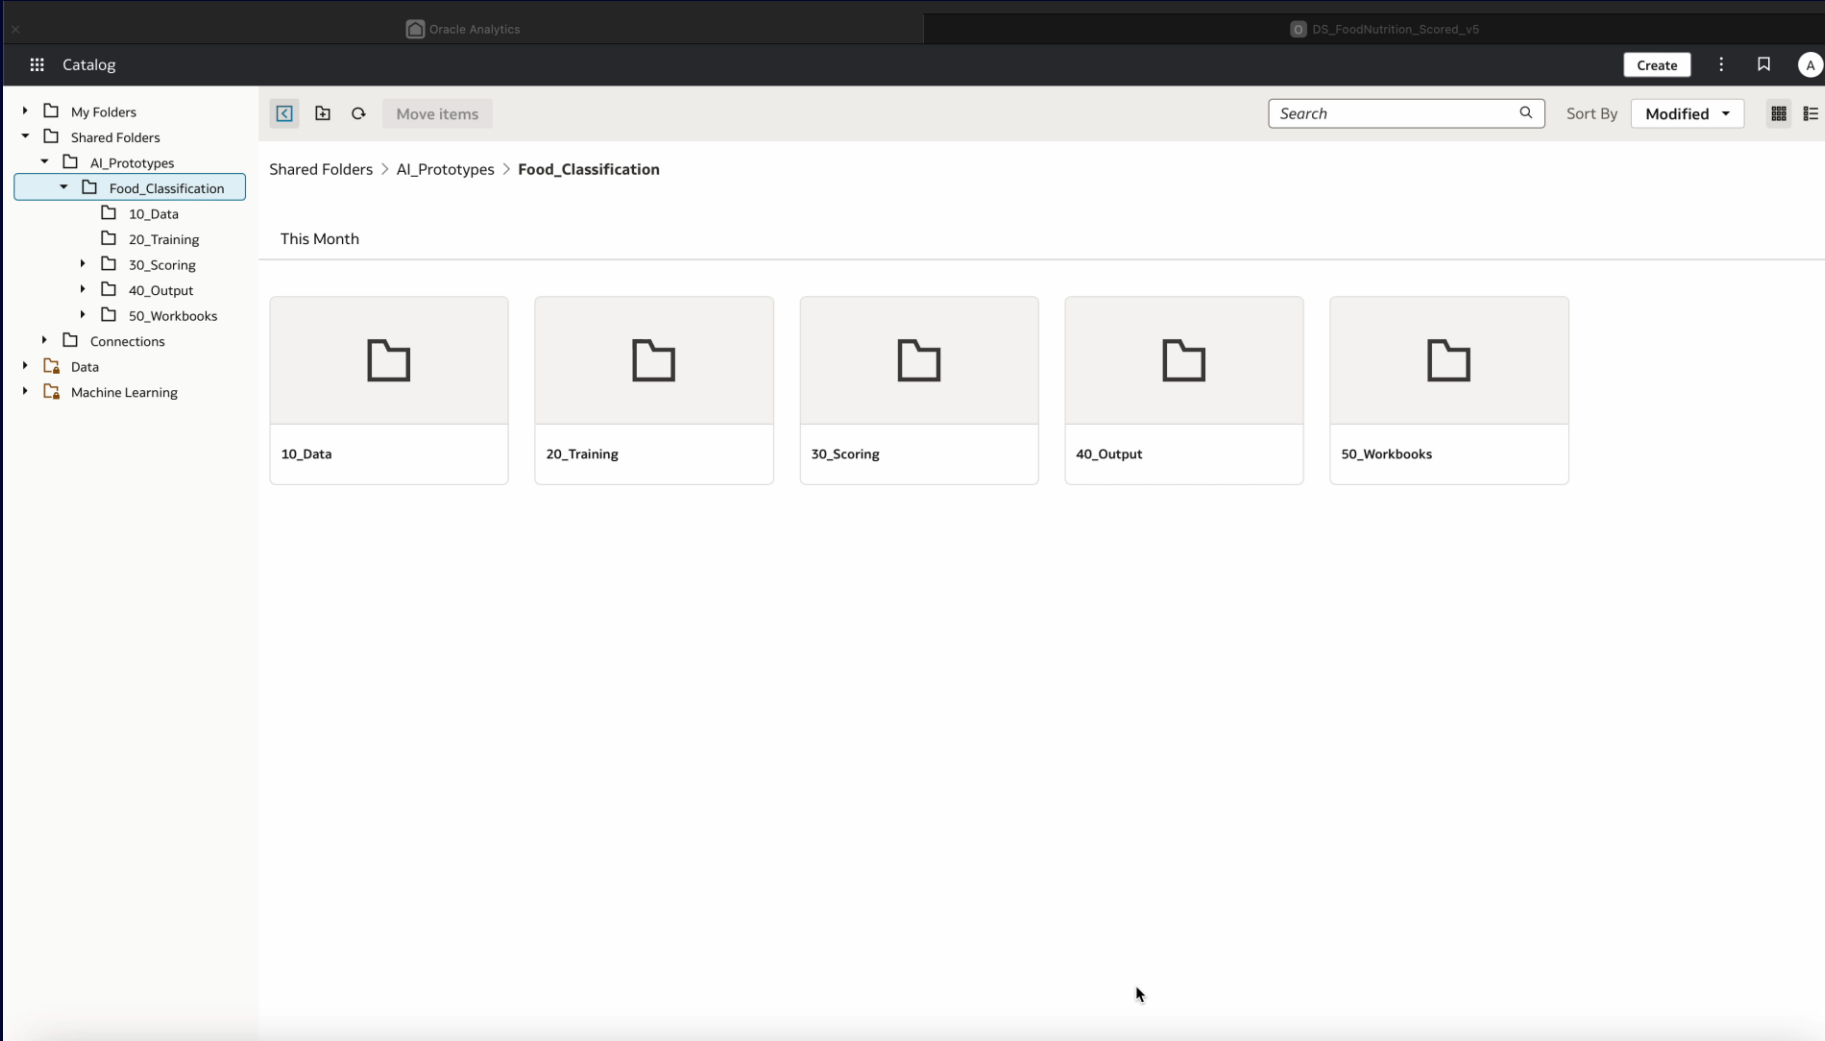
Accessing the Created Dataset

Opening the prepared dataset that will power our AI model



Creating our Model

Selecting the target variable and generating a machine learning model



The screenshot displays the Oracle Analytics Catalog interface. The left sidebar shows a tree view of folders: My Folders, Shared Folders, AI_Prototypes (expanded), Food_Classification (expanded), 10_Data, 20_Training, 30_Scoring, 40_Output, 50_Workbooks, Connections, Data, and Machine Learning. The main content area shows the path Shared Folders > AI_Prototypes > Food_Classification. Below the path, there is a 'This Month' section with five folder icons labeled 10_Data, 20_Training, 30_Scoring, 40_Output, and 50_Workbooks. The interface includes a search bar, a 'Move items' button, and a 'Sort By Modified' dropdown menu.



Random Forest Modeling

ML method that combines many decision trees to improve prediction accuracy

How it works

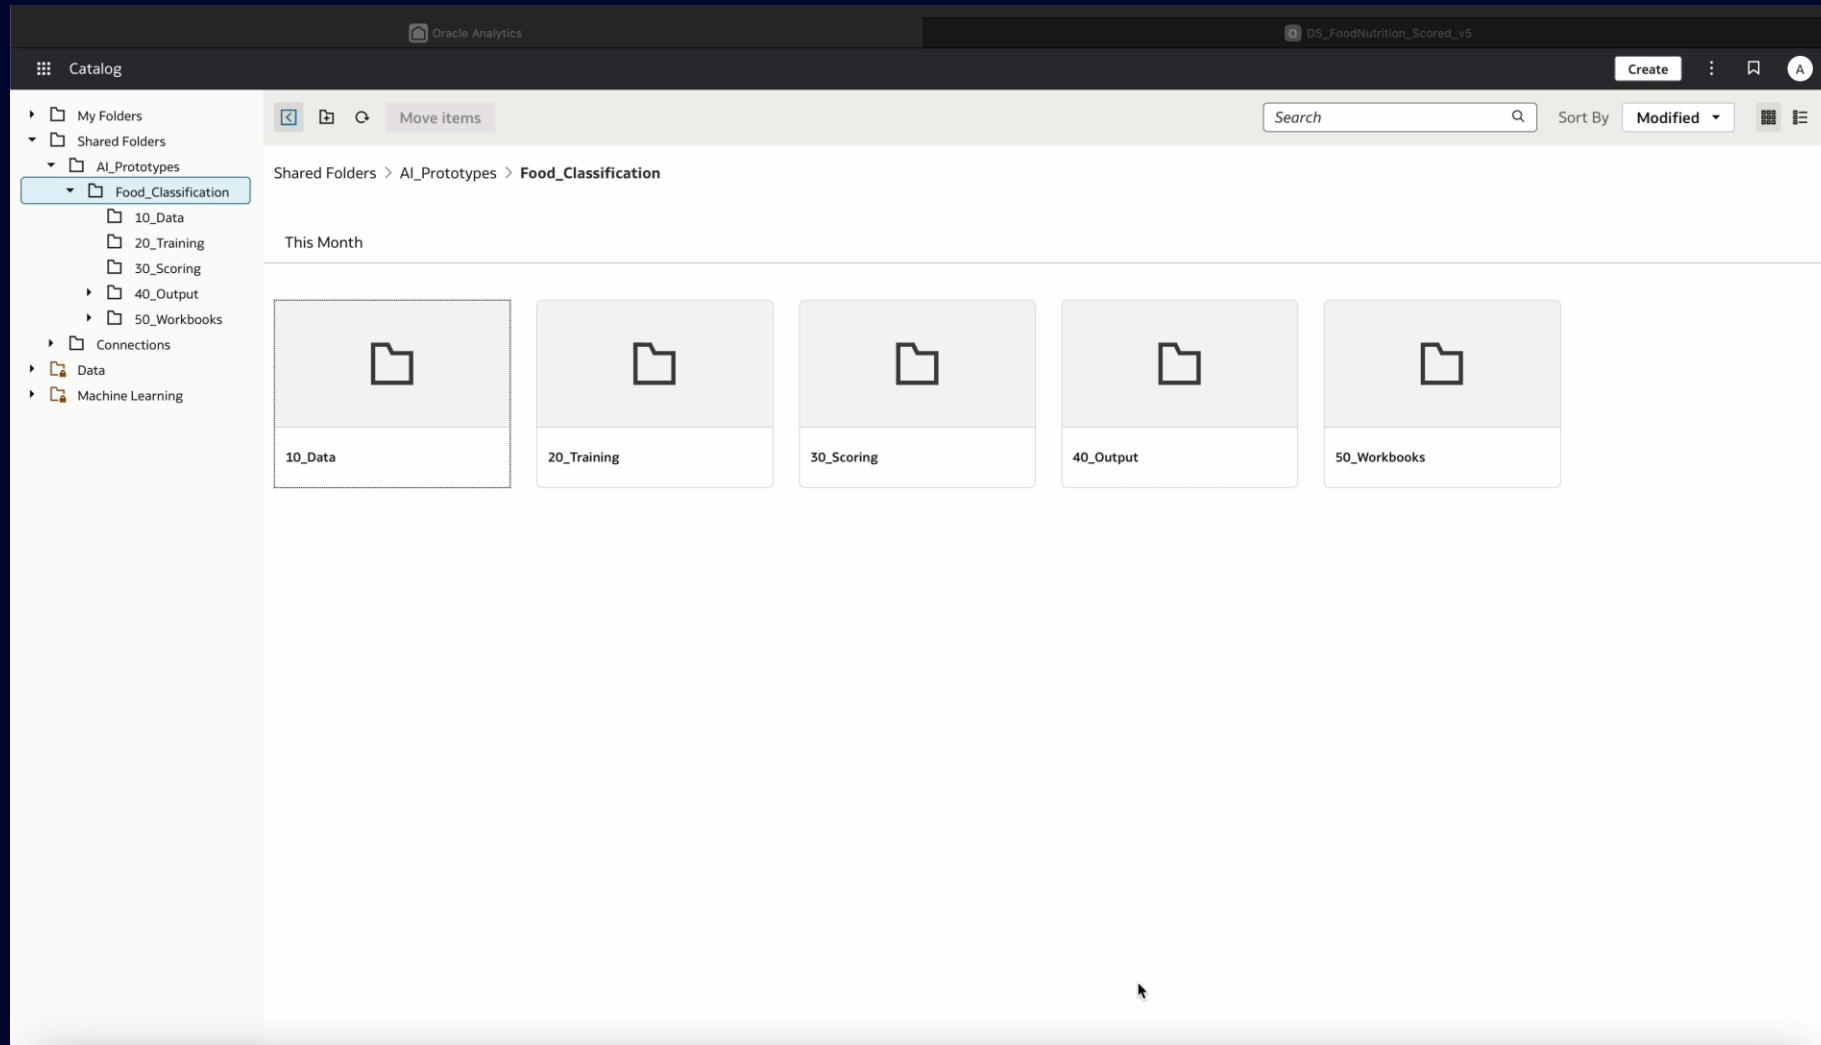
- Creates many decision trees
- Each tree makes a prediction
- Combines the predictions

Parameters

- Number of trees
- Sample size for a tree
- Number of features for a tree
- Minimum node size
- Balancing method

Applying our Model

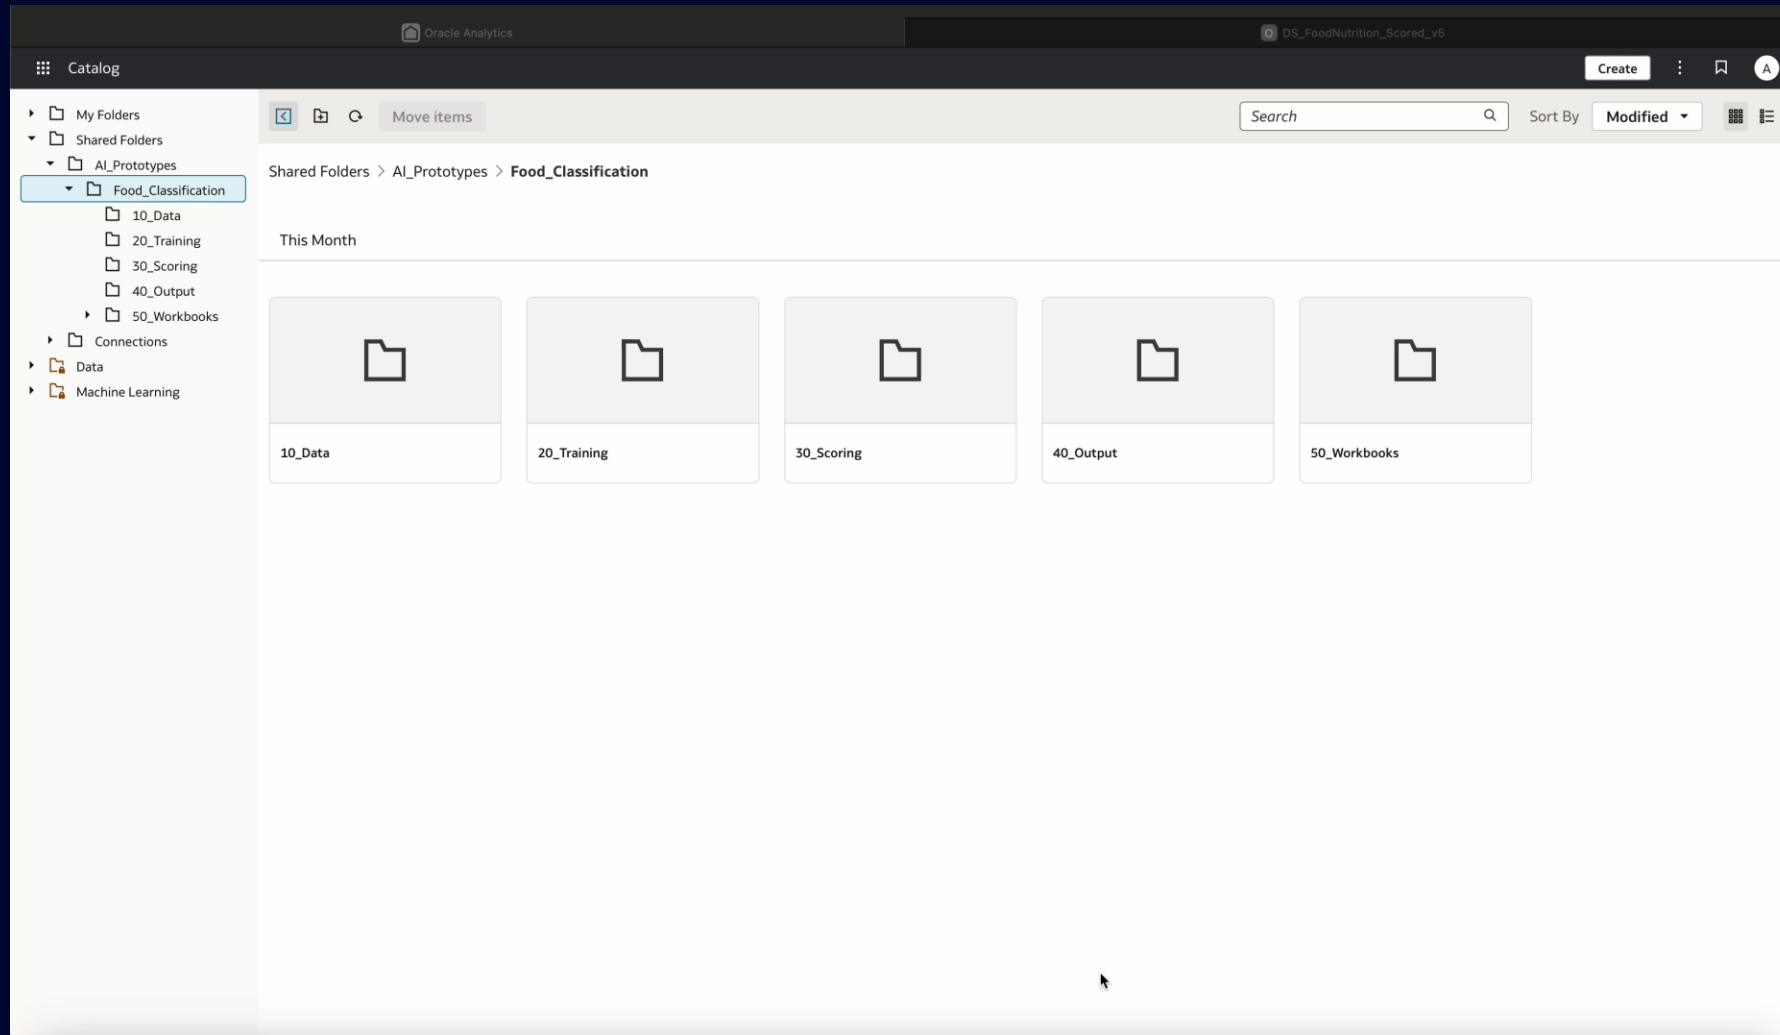
Running the model to generate predictions on our data



The screenshot displays the Oracle Analytics interface. The top navigation bar includes 'Catalog', 'Create', and user profile icons. The left sidebar shows a tree view with 'My Folders', 'Shared Folders', 'AI_Prototypes', 'Food_Classification', 'Connections', 'Data', and 'Machine Learning'. The main content area shows the path 'Shared Folders > AI_Prototypes > Food_Classification' and a 'This Month' view containing five folder icons labeled '10_Data', '20_Training', '30_Scoring', '40_Output', and '50_Workbooks'. The interface also features a search bar, a 'Sort By' dropdown set to 'Modified', and a 'Move items' button.

Viewing the Output

Interpreting predictions and exploring the model results



The screenshot displays the Oracle Analytics Catalog interface. The top navigation bar includes 'Catalog', 'Create', and user profile icons. The left sidebar shows a tree view with 'My Folders', 'Shared Folders', 'AI_Prototypes', 'Food_Classification', 'Connections', 'Data', and 'Machine Learning'. The main content area shows the breadcrumb path 'Shared Folders > AI_Prototypes > Food_Classification' and a 'This Month' filter. Below the filter, five folder icons are displayed, each with a label: '10_Data', '20_Training', '30_Scoring', '40_Output', and '50_Workbooks'. The interface also features a search bar, a 'Move items' button, and a 'Sort By Modified' dropdown menu.

AI Prototyping Made Practical

Build models, generate predictions, and view output

Workflow summary:

- 1 Access the dataset
 - Start with curated business data in Oracle Analytics
- 2 Create the model
 - Select the target outcome and generate a ML model
- 3 Apply the model
 - Run the model against your data for predictions
- 4 Review the output
 - Analyze results, patterns, and predicted outcomes
- 5 Turn insights into action
 - Use the output to support decisions, prototype use cases, and guide next steps

05

Data Science

While Oracle Analytics Cloud makes machine learning accessible, Oracle Data Science lets us go deeper.

Unit Testing Code & Repository

Explore and Use It!






 GitHub Repository

This repository contains the scripts used to validate an Oracle Data Science notebook environment before building AI models.



https://github.com/DennisFromDefinian/data_science_notebook_unit_test

What these scripts verify

-  Runtime environment readiness
-  OCI authentication (Resource Principals)
-  Object Storage connectivity
-  Bucket permissions
-  End-to-end data read/write operations





Goal

Why Unit Tests for AI Infrastructure?

Reliable prototypes start with reliable environments

AI experiments often fail because of **environment issues**, not models.

Common causes:

-  Networking or DNS issues
-  IAM policy misconfiguration
-  Storage access failures
-  Runtime environment differences

Unit tests help us

- ✓ Validate infrastructure quickly
- ✓ Detect configuration issues early
- ✓ Confirm access to cloud services

Result:





More time spent building AI — less time debugging infrastructure.

Environment Sanity Checks





Is the notebook runtime ready to run AI workloads?

Script: `01_env_sanity_check.py`

Tests basic platform functionality:

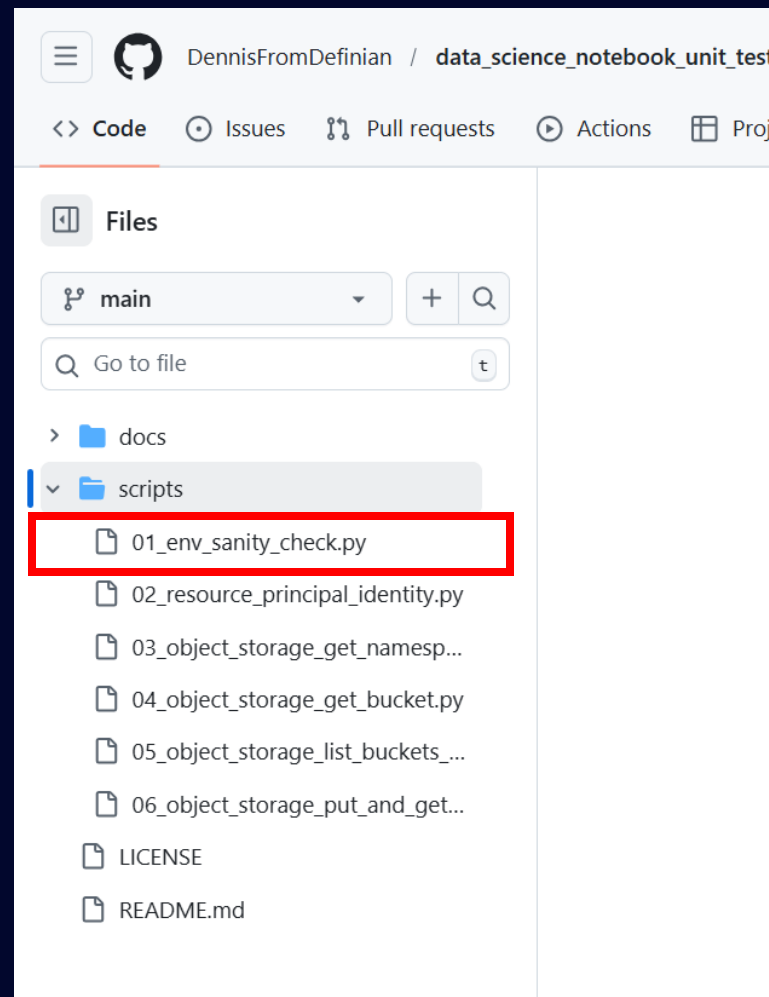
-  Python runtime verification
-  DNS resolution checks
-  HTTPS connectivity tests
-  Environment variable validation

Libraries used:

-  `os`
-  `sys`
-  `socket`
-  `urllib`

Why this matters:

Infrastructure misconfiguration is one of the **most common causes of prototype failure.**



Secure Authentication with OCI

Enterprise AI without hard-coded credentials




Script: `02_resource_principal_identity.py`

Uses the **OCI Python SDK**

Key function:

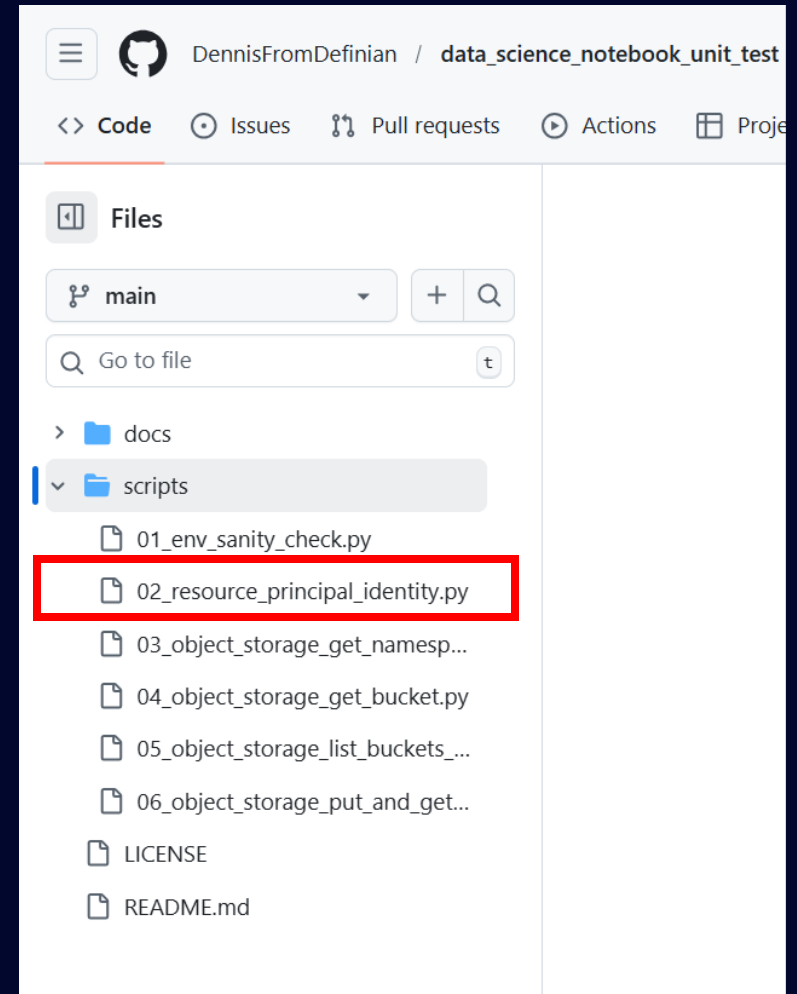
```
oci.auth.signers.get_resource_principals_signer()
```

Validates:

-  Notebook identity
-  Tenancy access
-  IAM policy configuration

Why it matters:

Enterprise environments require **secure authentication without embedded credentials.**



Testing Object Storage Access

Validating data access for AI pipelines




Scripts:

- 03_object_storage_get_namespace.py
- 04_object_storage_get_bucket.py
- 05_list_buckets.py

OCI Service:

 **Object Storage**

These tests verify:

-  Bucket visibility
-  IAM permissions
-  Region configuration

Why this matters:

Most AI workflows depend on **object storage for training data and artifacts.**



End-to-End Data Pipeline Test

Proving the environment can move data

06_object_storage_put_and_get_object.py

Workflow:

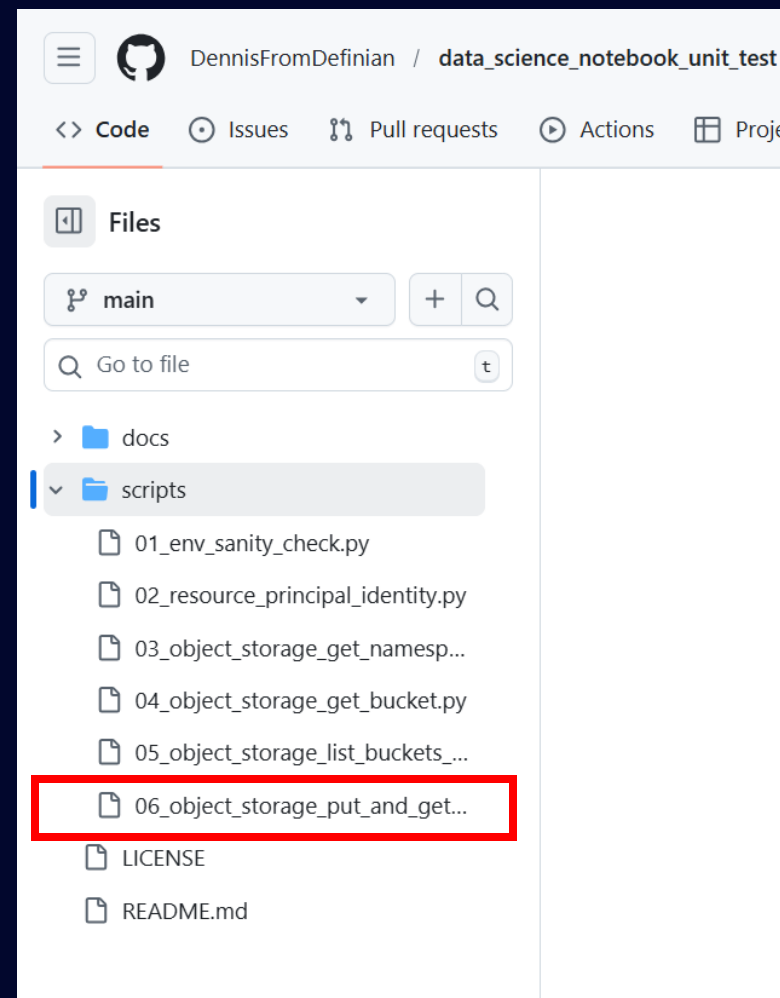
- 1 Write JSON data to Object Storage
- 2 Retrieve the object
- 3 Validate the returned data

Libraries used:

-  oci
-  json
-  datetime

What this proves:

- ✓ IAM permissions
- ✓ Storage connectivity
- ✓ Data pipeline readiness



Prototype Code & Repository

A practical Oracle Data Science workflow you can explore after the session

This prototype demonstrates:

- 🏠 API-based open data ingestion
- 🖌️ text normalization and preparation
- 🧠 embedding generation
- 🔍 FAISS vector search
- ☁️ Object Storage persistence
- 📖 table + map visualization



https://github.com/DennisFromDefinian/data_science_notebook_prototype

Shows how a **small team** can build an end-to-end AI prototype in **OCI Data Science** using public data and cloud-native components.

What These Scripts Are Trying to Achieve

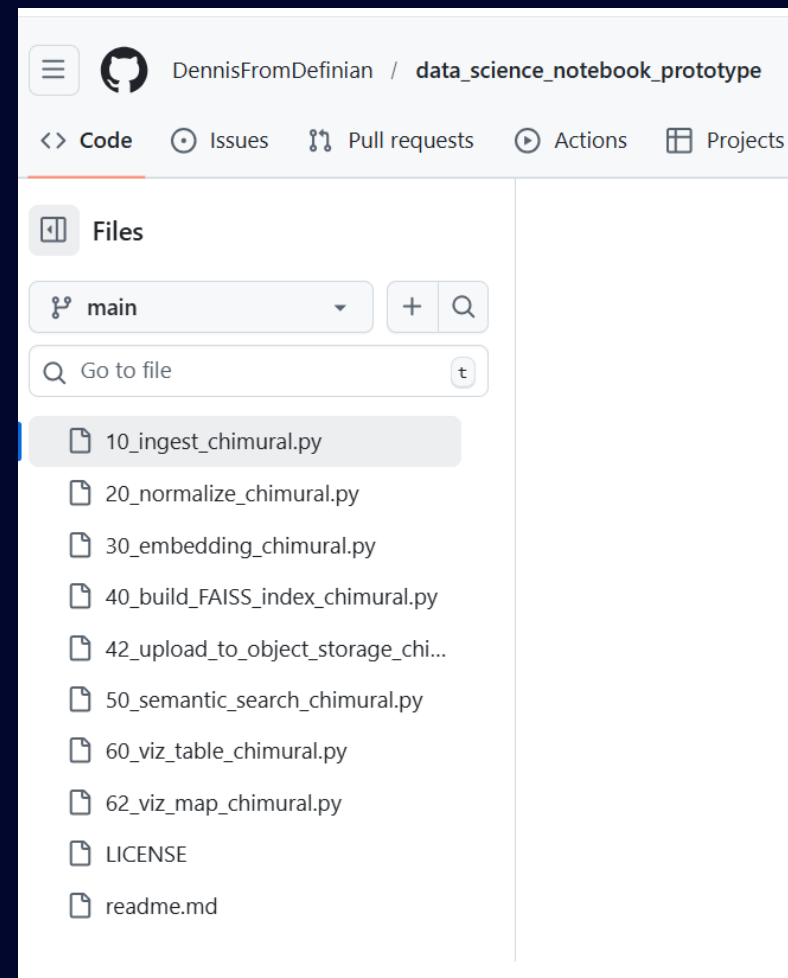
From raw public data to searchable, usable insight

Pipeline flow

- 1 Ingest mural data from a public API
- 2 Normalize and curate the text
- 3 Turn text into vector embeddings
- 4 Build a searchable FAISS index
- 5 Persist artifacts to Object Storage
- 6 Run semantic search
- 7 Present the results in human-friendly form

Why this matters






AI prototypes create value when they turn **messy data into retrievable business insight.**






Tools, Libraries, and OCI Building Blocks

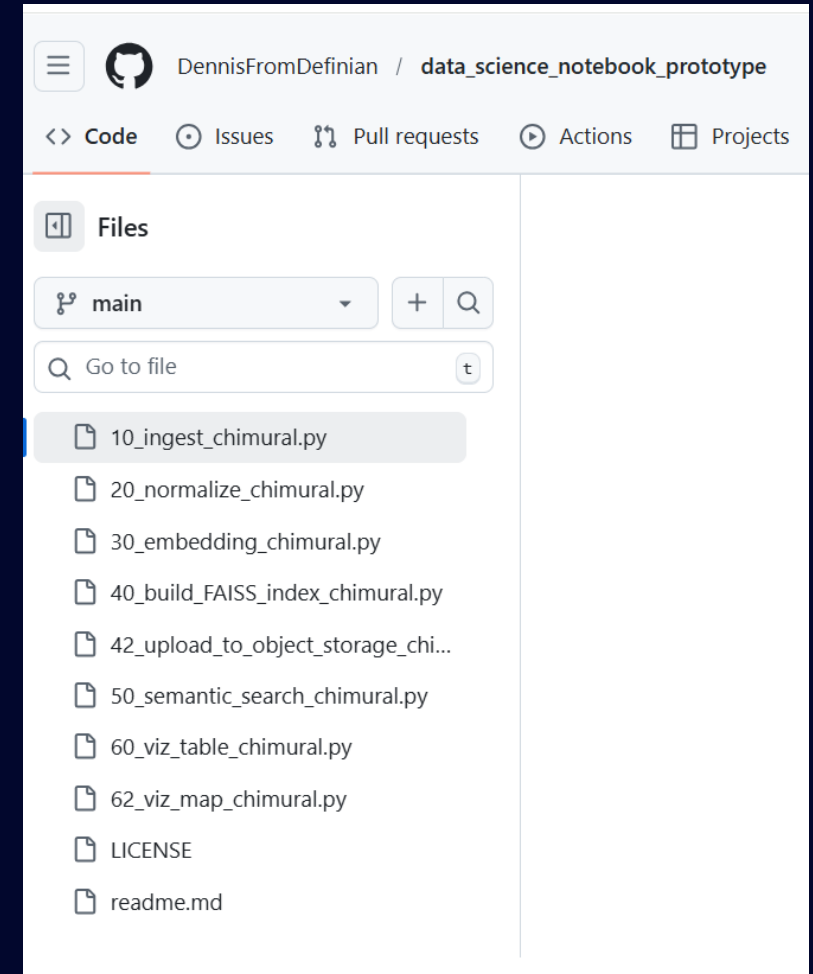
Simple components, combined intentionally

Python / notebook tools

-  requests + pandas for ingestion and shaping
-  embedding models via transformer-style workflow
-  faiss for vector similarity search
-  oci SDK + Resource Principals for Object Storage
-  folium for map-based visualization

OCI concepts

-  OCI Data Science notebook sessions
-  Object Storage for persistent artifacts
-  policy-driven access and cloud-native security



unit_testing.ipynb x Untitled.ipynb x

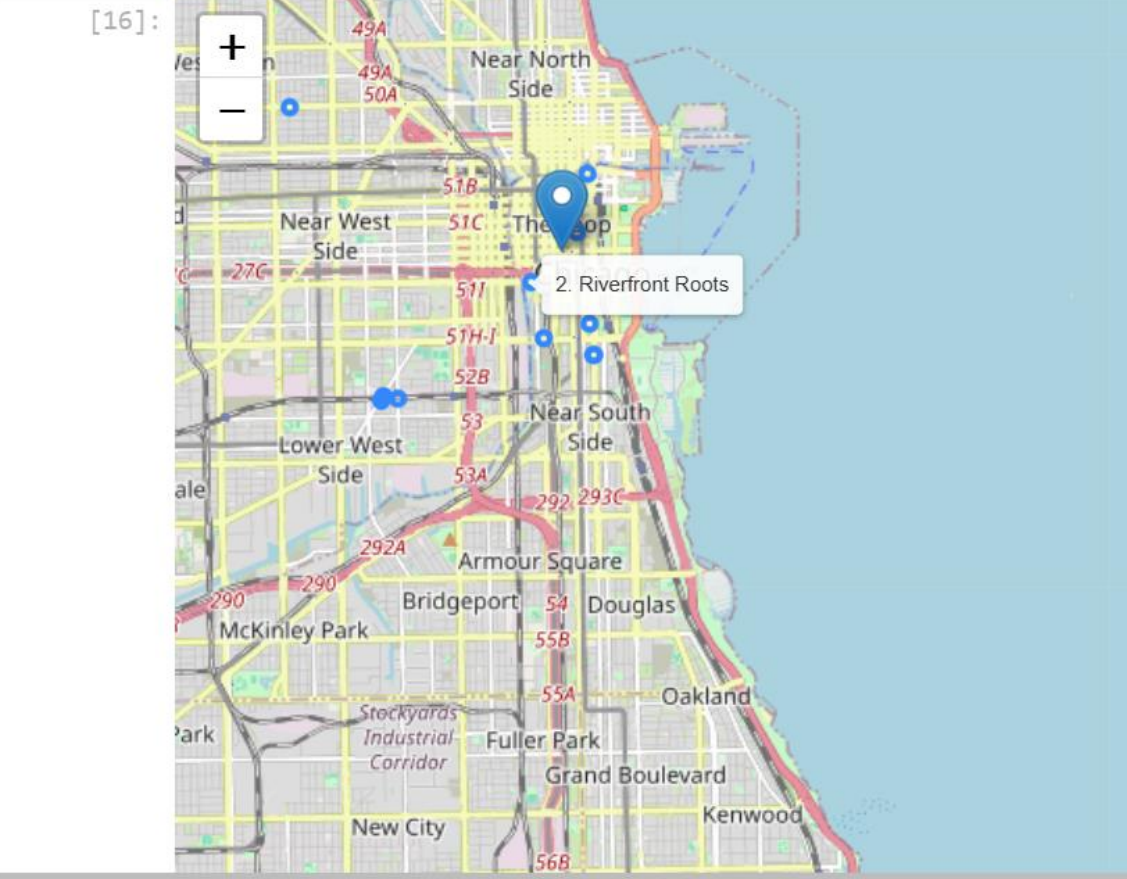
Python [conda env:root] *

```
# Note: We retrieve more results first (k=50) then filter down to those within radius.
semantic_search(
    "nature scene",
    k=50,
    center_lat=41.8781, # downtown Chicago-ish
    center_lon=-87.6298,
    radius_miles=3.0,
).head(5)
```

[13]:

	rank	query	score	distance_miles	mural_id	title	artist	community_area	latitude	longitude	search_text
0	1	nature scene	0.459246	0.696992	19287	untitled	Revise CMW	NaN	41.887615	-87.625301	untitled Nature found amidst the chaos and gri...
1	2	nature scene	0.381207	0.913538	19205	Wild Life	Eelco Van Den Berg	33	41.865539	-87.624257	Wild Life (Photo Credit: Sandra Steinbrecher) ...


Name	Modified
lost+found	18d ago
chicago_murals_m...	11d ago
chicago_murals.fai...	11d ago
unit_testing.ipynb	14d ago
Untitled.ipynb	11d ago





Business Concepts This Prototype Maps To


Murals are the example, retrieval is the real lesson

Concrete business parallels

 Manufacturing: search maintenance notes for “vibration and overheating”

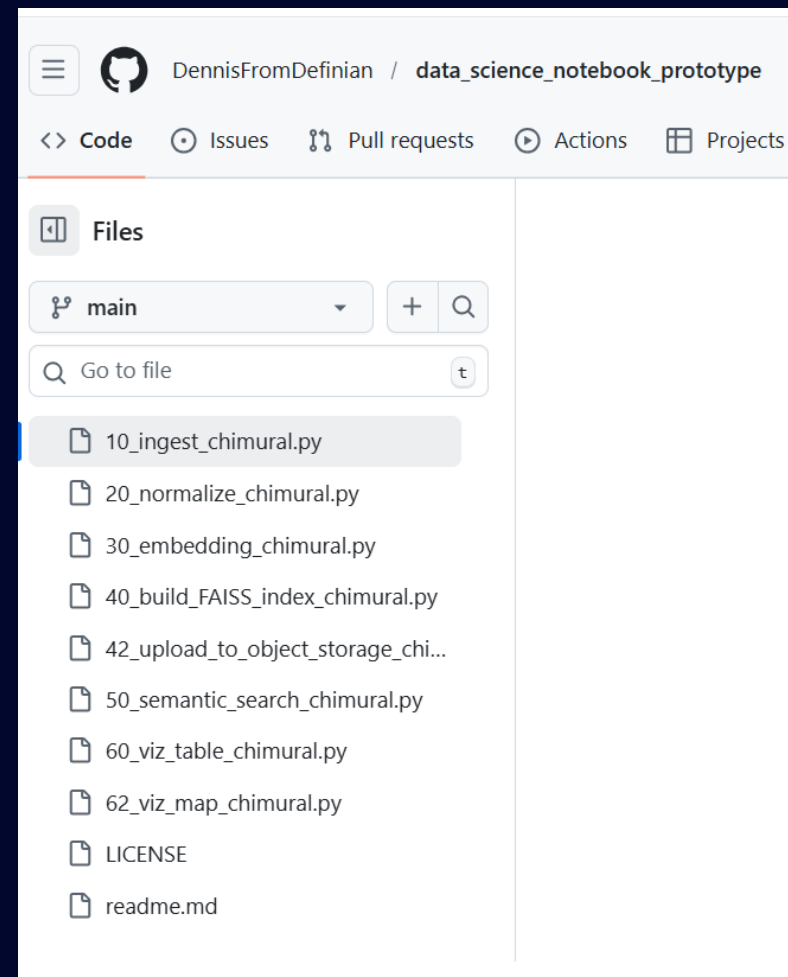
 Supply chain: find similar supplier issues from free-text comments

 Service operations: retrieve tickets that “sound like” the current issue

 Product data: search items by descriptive similarity, not just exact keywords

 Field ops: combine semantic retrieval with location context


Better retrieval from unstructured text can improve **triage, discovery, reuse, and decision-making.**




This Flow Is Not Canonical

There are several valid ways to build this


You might choose differently:

 **Data landing:** CSV in Object Storage instead of live API calls

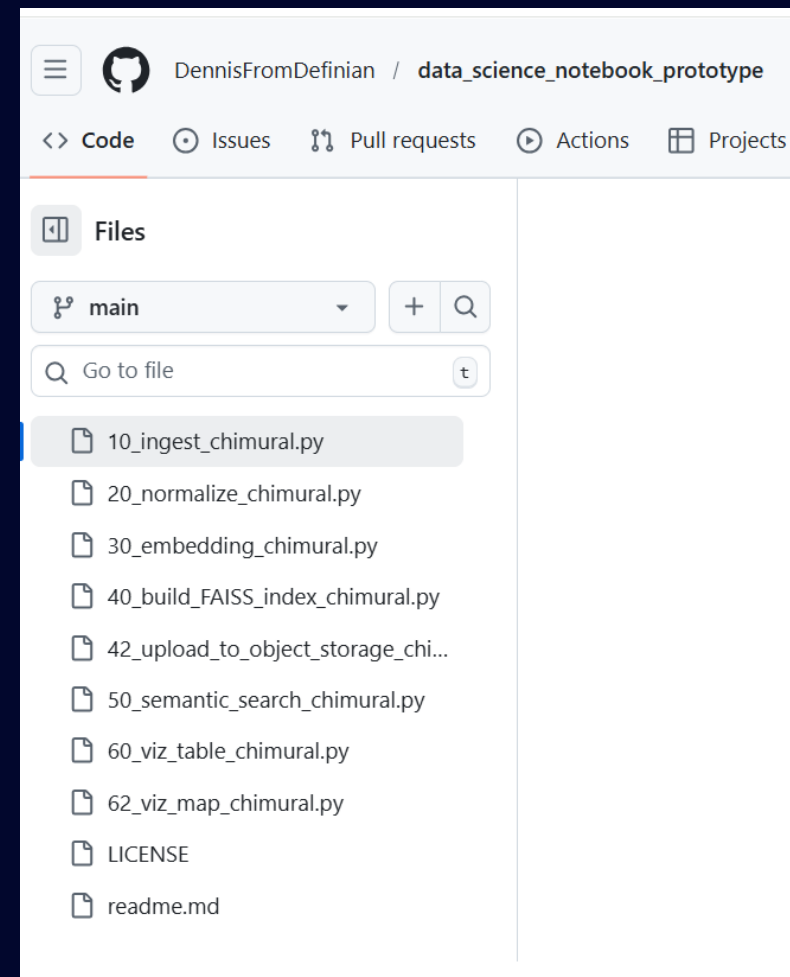
 **Text prep:** different field combinations or more aggressive cleaning

 **Embeddings:** another model, hosted endpoint, or database-native option

 **Indexing:** Oracle vector features instead of FAISS

 **Orchestration:** notebooks, scripts, pipelines, or application services






 **Presentation:** OAC dashboards, apps, or APIs instead of notebook output



What This Says About Small-Team AI Prototyping

Practical, modular, & close to real enterprise needs

This repo reflects design principles such as:

-  modular execution
-  clear pipeline staging
-  cloud-native artifact management
-  low-friction experimentation with public data
-  small-team feasibility

And from here...

I want to close this technical section by briefly introducing the company I work for, **Definian** — because this mindset of turning complex data problems into practical, usable solutions is central to what we do.



06

Definian

A little about ourselves

Defining What's Next with Data

Definian is the enterprise data transformation partner trusted by prominent brands to turn complexity into clarity, risk into readiness, and data into business value.



40+
years in business

Legacy of enterprise data expertise, now powered for the AI era.

Over **1,000**
successful
projects

Trusted by leading brands to tackle the most complex data challenges.

Clientele of
Fortune 500 &
Top SIs

A full-spectrum data partner
Strategy, Modernization,
Insights & AI.

The Speakers Today



Dennis Gray

<https://www.linkedin.com/in/dennisgray>

<https://github.com/dennisfromdefinian>



Adam Hume

<https://www.linkedin.com/in/adam-hume-a169aa1b5/>



Thank You

Helpful Links –

ORACLE ANALYTICS VIDEOS:

<https://www.youtube.com/@OracleAnalytics/videos>

OAC NOVEMBER 2025 AND JANUARY 2026 NEW FEATURES VIDEOS BY ORACLE:

<https://bit.ly/OACNov25Features>

<https://bit.ly/OACJan26Features>

OAC NEW FEATURES DOCUMENTATION BY ORACLE:

<https://docs.oracle.com/en/cloud/paas/analytics-cloud/acswm/index.html#GUID-CFF90F44-BCEB-49EE-B40B-8D040F02D476>

ORACLE ANALYTICS COMMUNITY:

<https://community.oracle.com/products/oracleanalytics>

ORACLE ANALYTICS LIBRARY/EXAMPLES:

<https://www.oracle.com/business-analytics/data-visualization/examples/>

ORACLE ANALYTICS LIVE DEMOS:

<https://www.oracle.com/business-analytics/data-visualization/demos/>

Past & Future TechCasts:



February 5th

Our Favorite Features of OAC:
November 2025 & January 2026 Releases

Presented by Tim Vlamis, Branden Pavol, Wayne Van Sluys, Gautam Pisharam & Taiwo Ajayi



March 5th

Fraud Stops with Oracle Database 26ai
Vector Search + OML

Presented by Abi Giles-Haigh



April 2nd

Build Practical AI Prototypes with Oracle
Analytics (Even with a Small Team)

Presented by Dennis Gray & Adam Hume

TechCast Archive

2026	2025	2024	2023	2022	2021
Date	Title	Presenter(s)		Replay	Download(s)
Mar 5	Fraud Stops with Oracle Database 26ai Vector Search + OML	Dr. Abigail Giles-Haigh		Video	Slides
Feb 5	Our Favorite Features of OAC: November 2025 & January 2026 Releases	Branden Pavol, Gautam Pisharam, & Carter Beaton		Video	Slides
Jan 8	The Oracle AI Microservices Sandbox for RAG Rapid Prototyping	Melli Annamalai & Phani Chilakapati		Video	Slides

Submit a topic to share at <https://andouc.org/techcasts/>

Feedback Panel

Oracle PM's:



Emcees:




- Provide feedback **directly** to Oracle PM's
- Find out what's **coming next**
- Tell them **what you'd like to see next** in future releases

Thursday - 1:30p



Analytics and Data

SUMMIT 2026



Only 1 week to go!
April 14-16, 2026



Scan this for full information:

