

#FABCONSQLCON2026

FABCON

Microsoft Fabric
COMMUNITY CONFERENCE

SQLCON

Microsoft SQL
COMMUNITY CONFERENCE

ATLANTA MARCH 16 - 20, 2026



Building MLOps for Traditional ML in Microsoft Fabric

Nicholas Leonhard

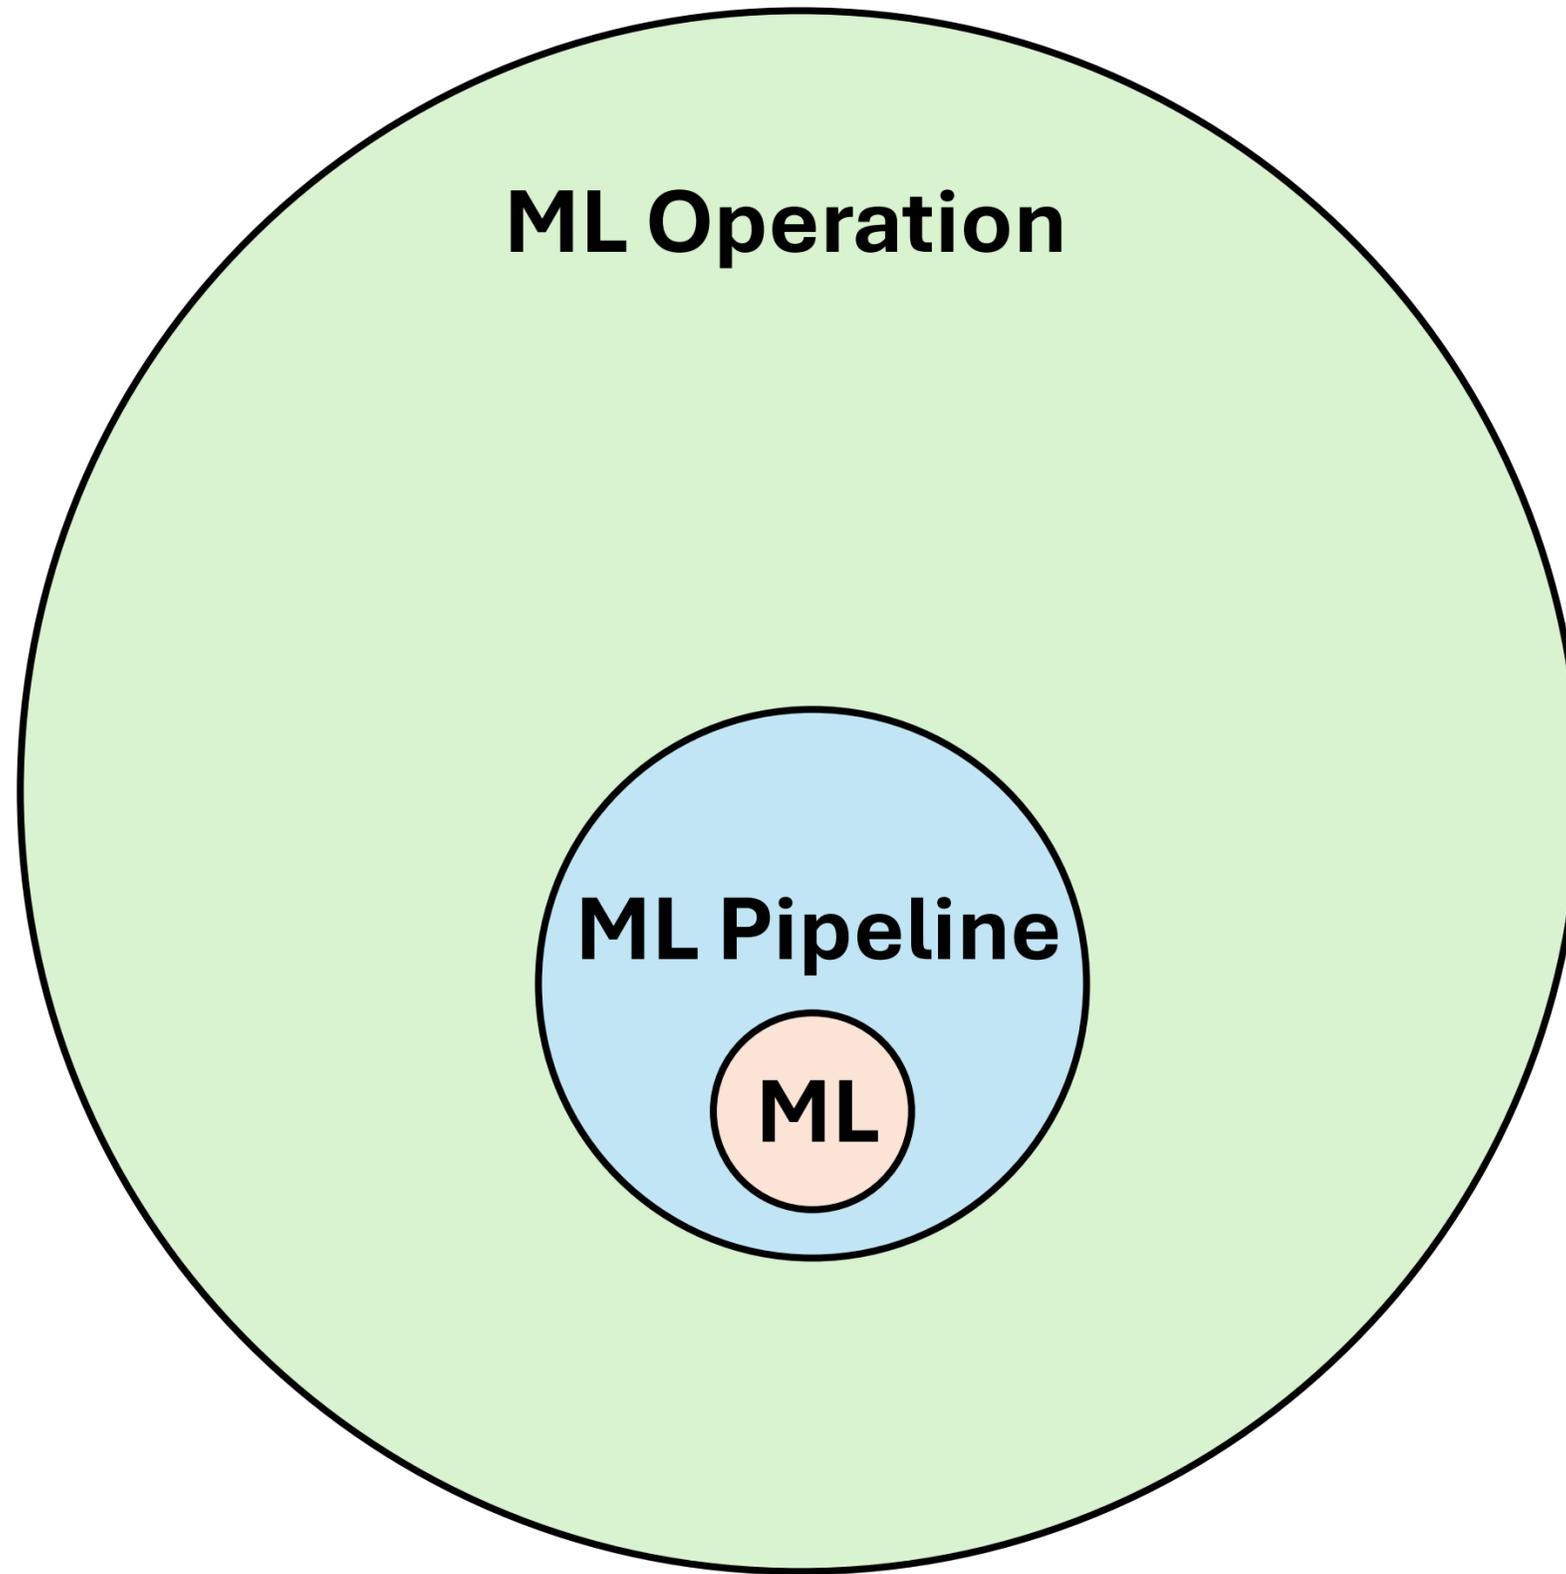
Machine Learning Engineer



Connect
on LinkedIn:

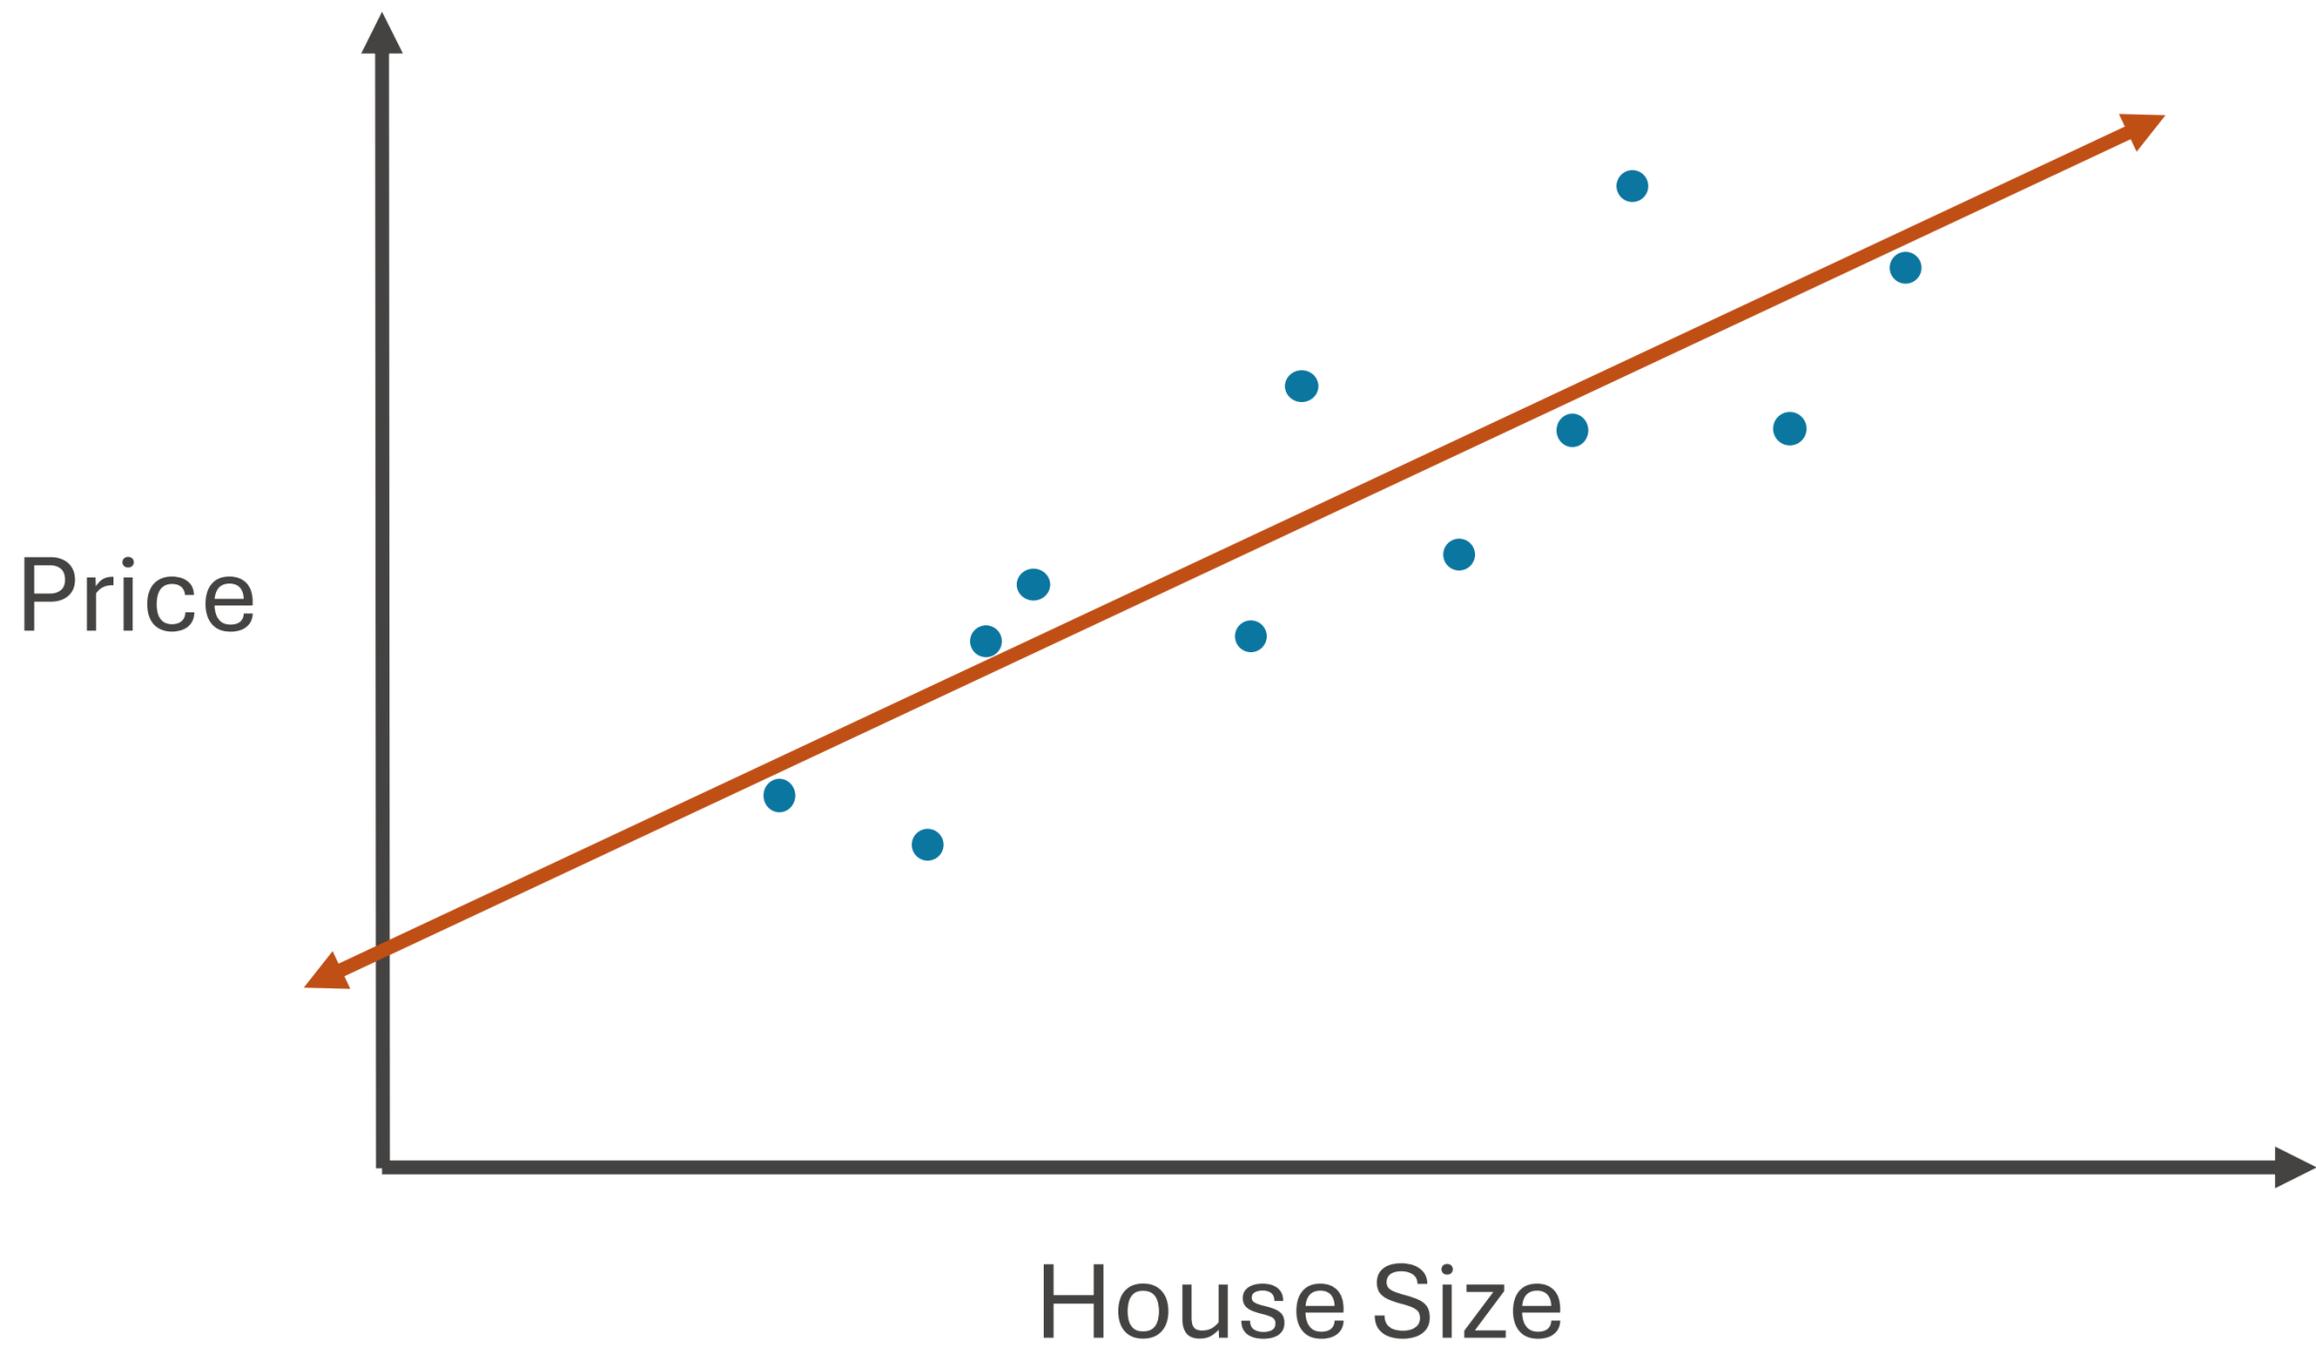
<https://www.linkedin.com/in/nicholas-leonhard-21a972ab>

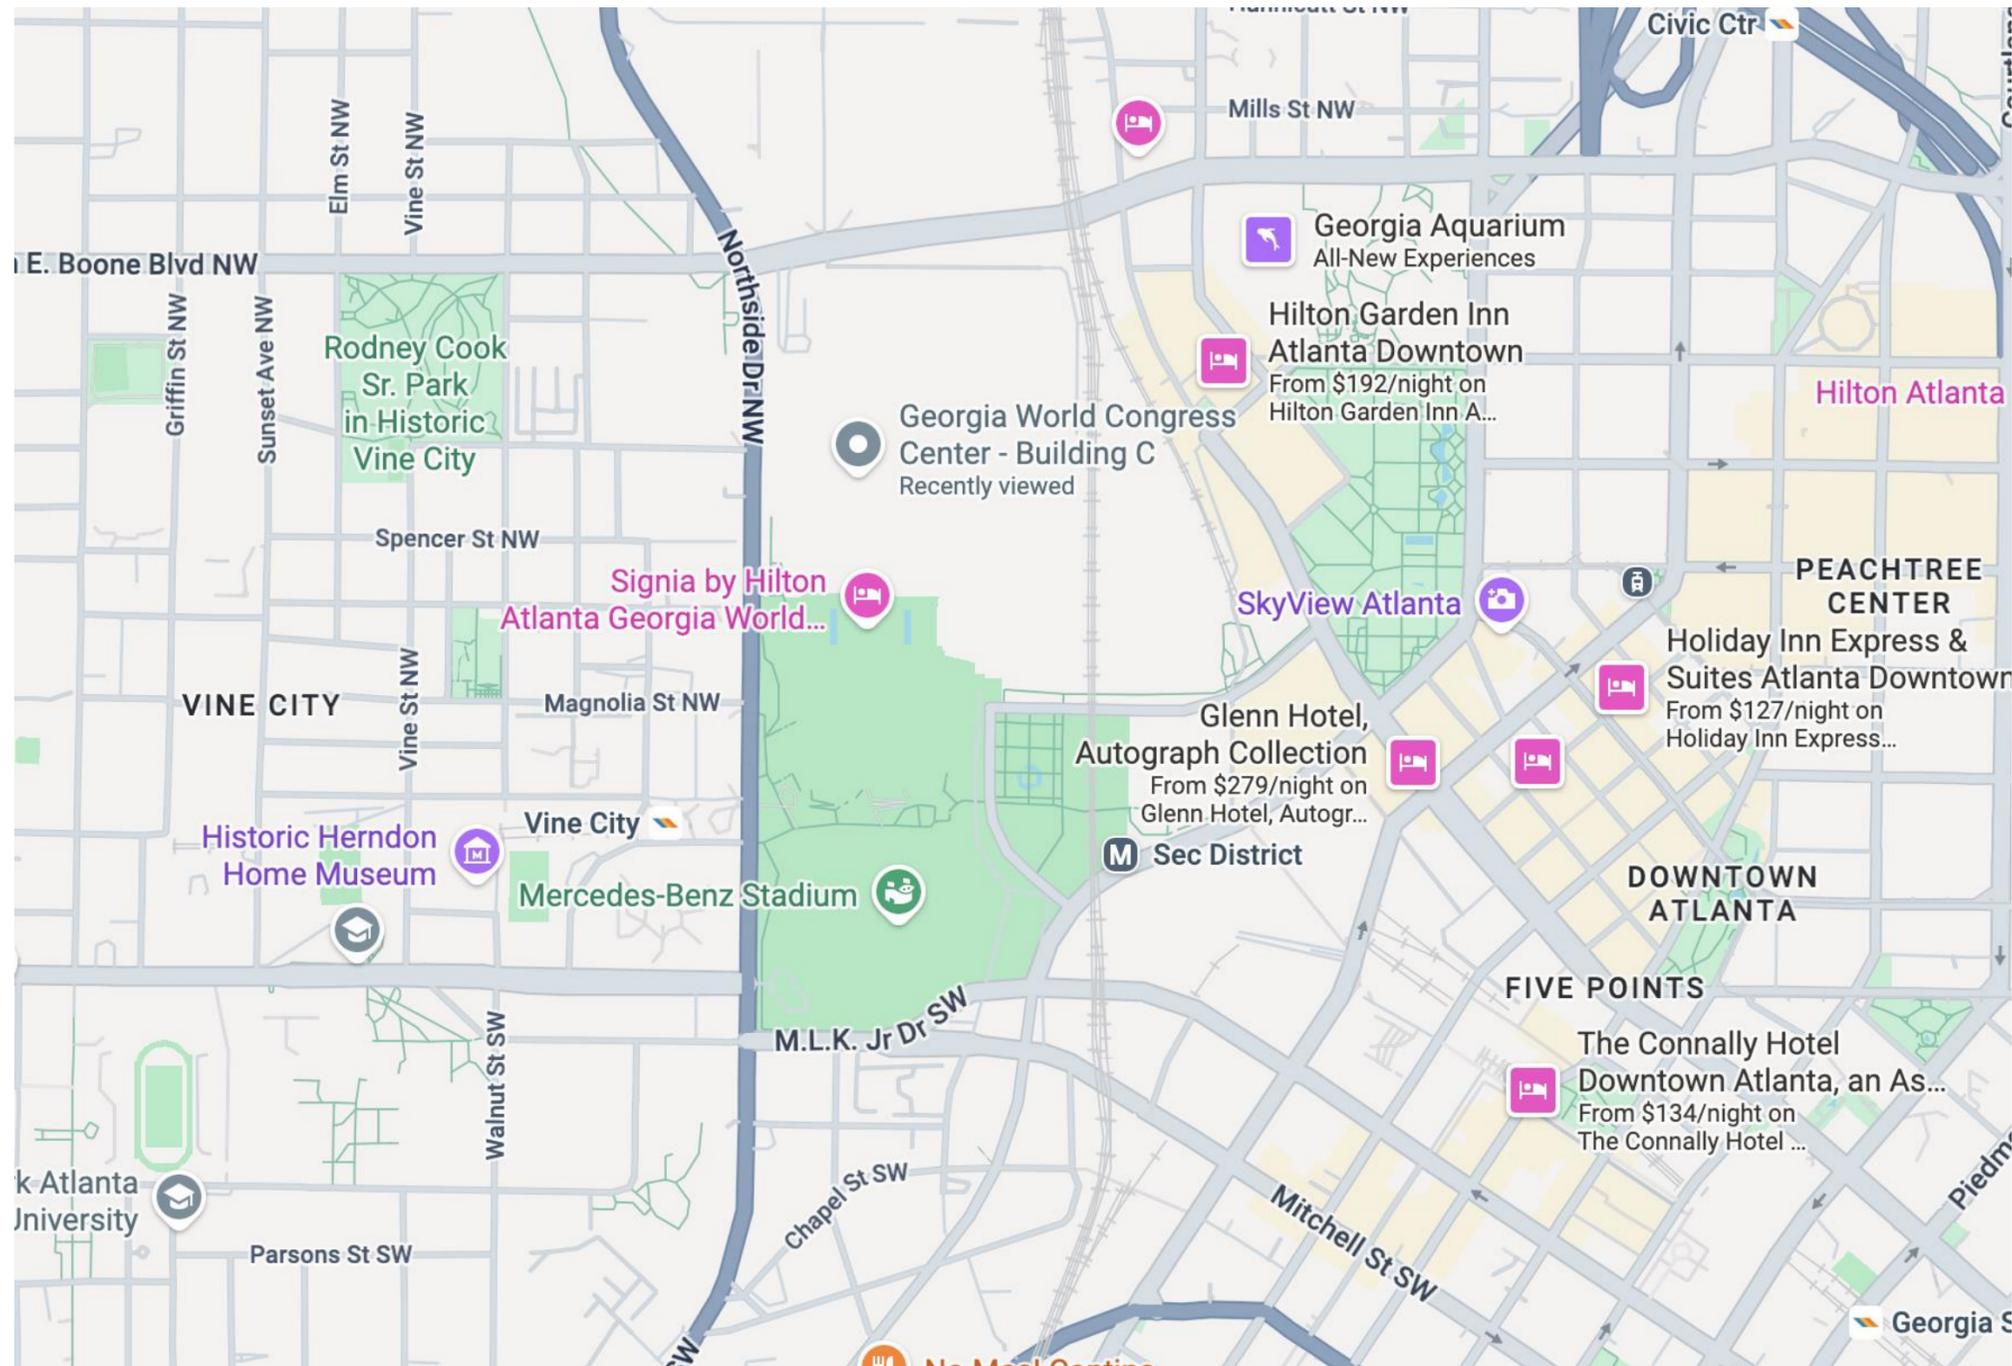
What is MLOps?

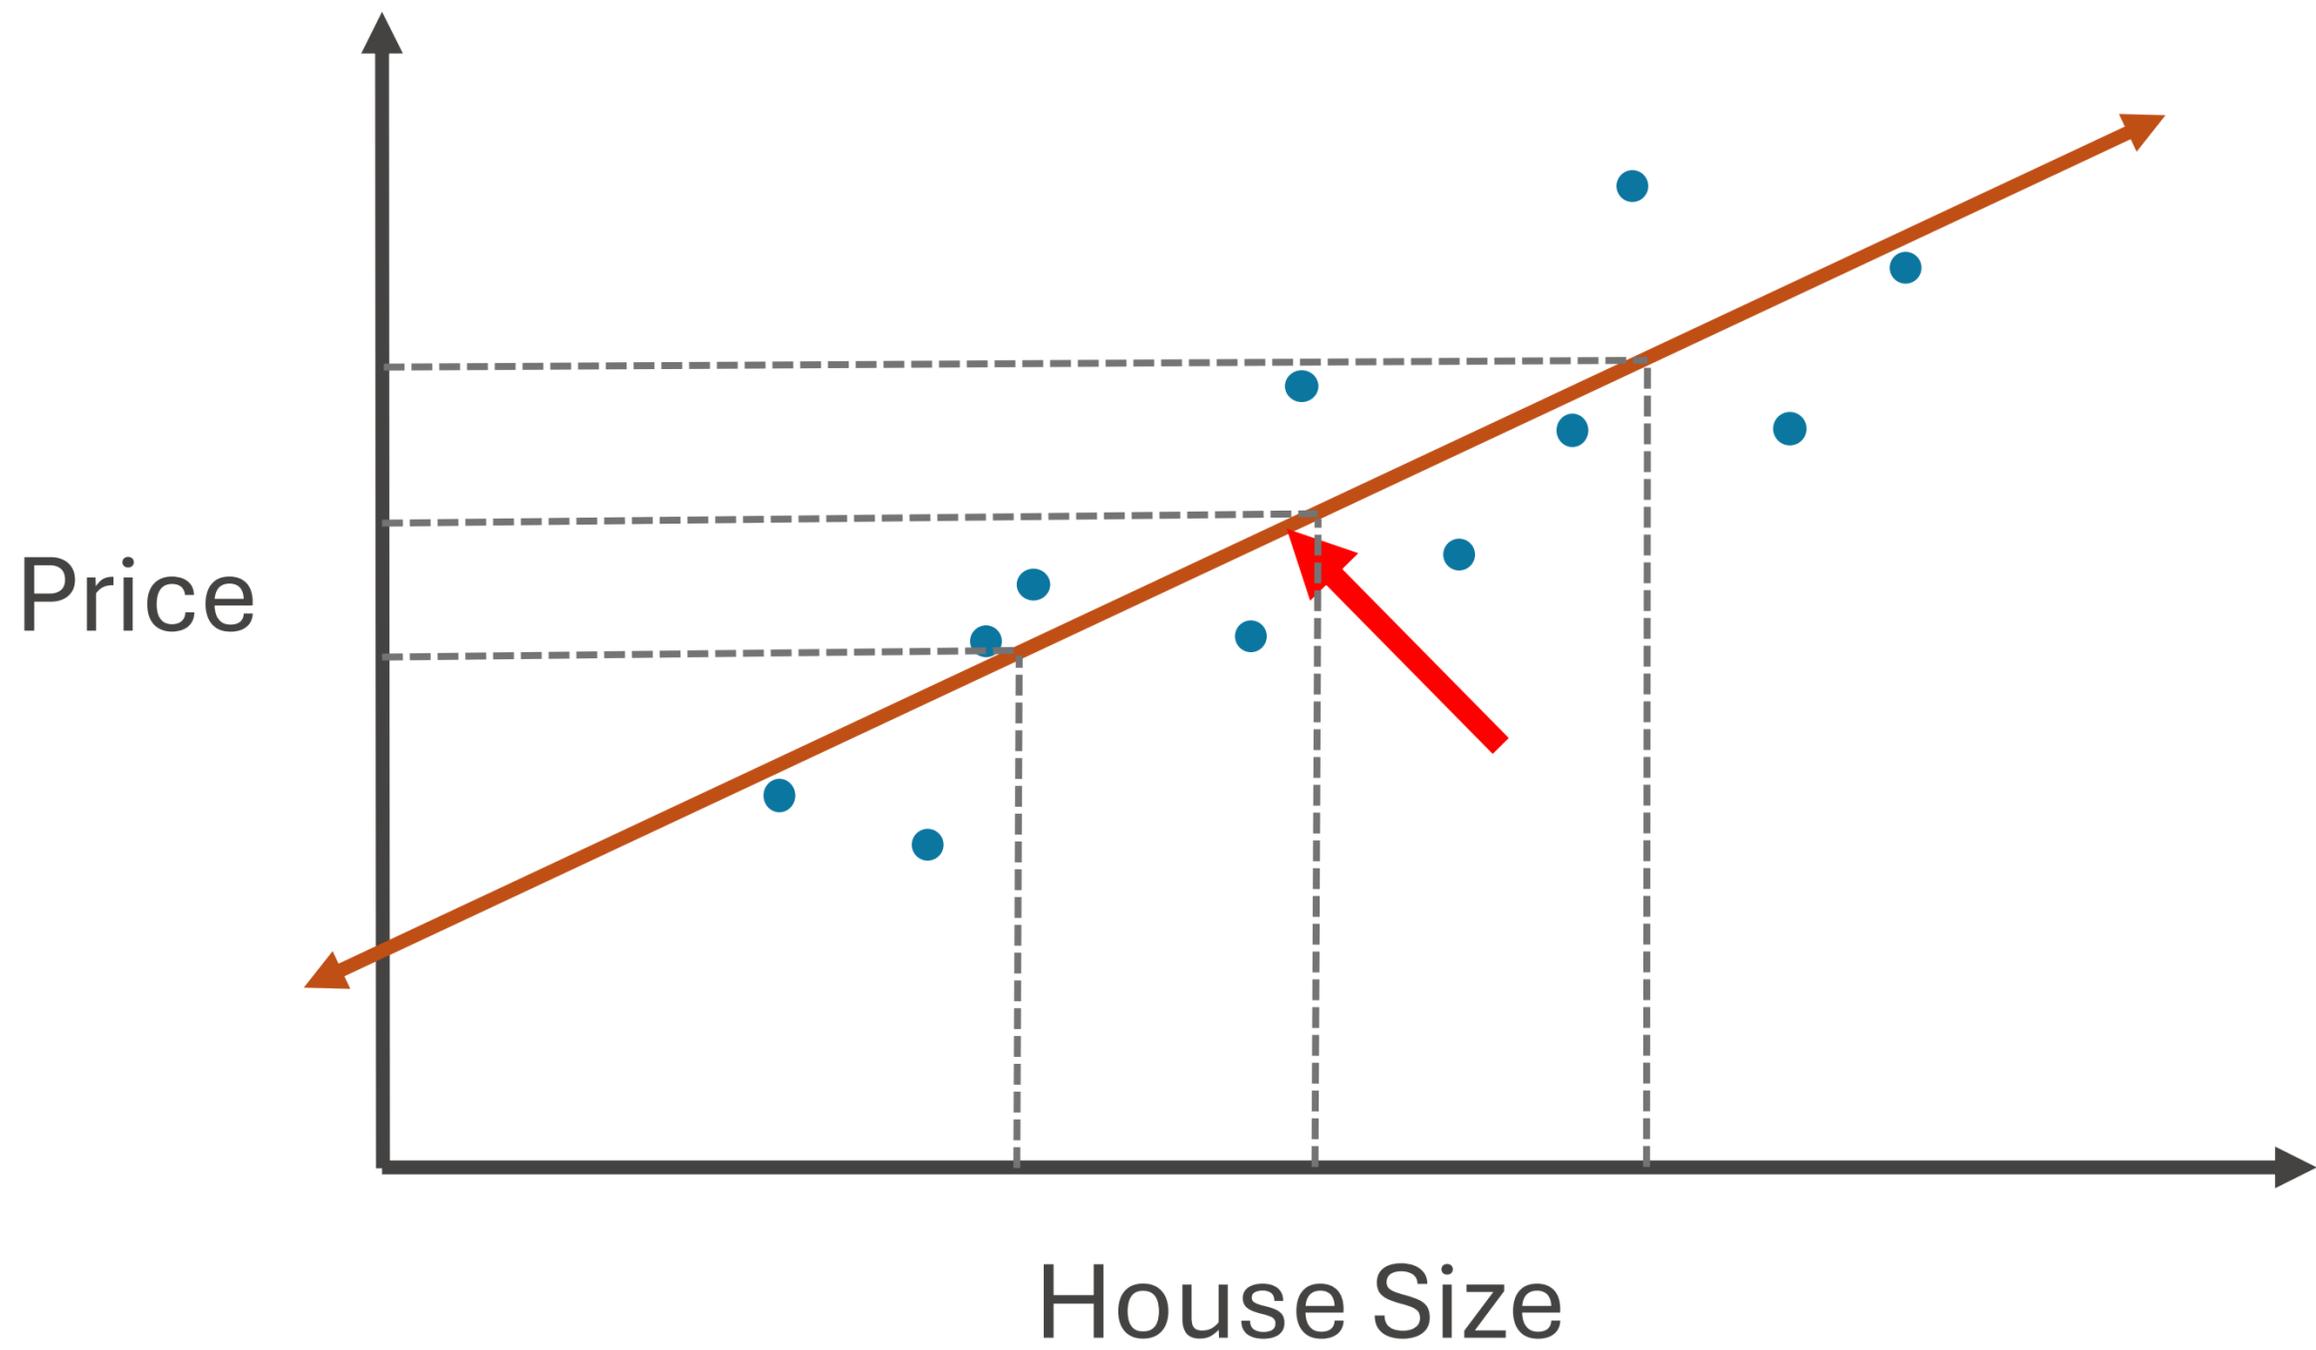


What is Machine Learning?

- Model a relationship







Statistical
Methods

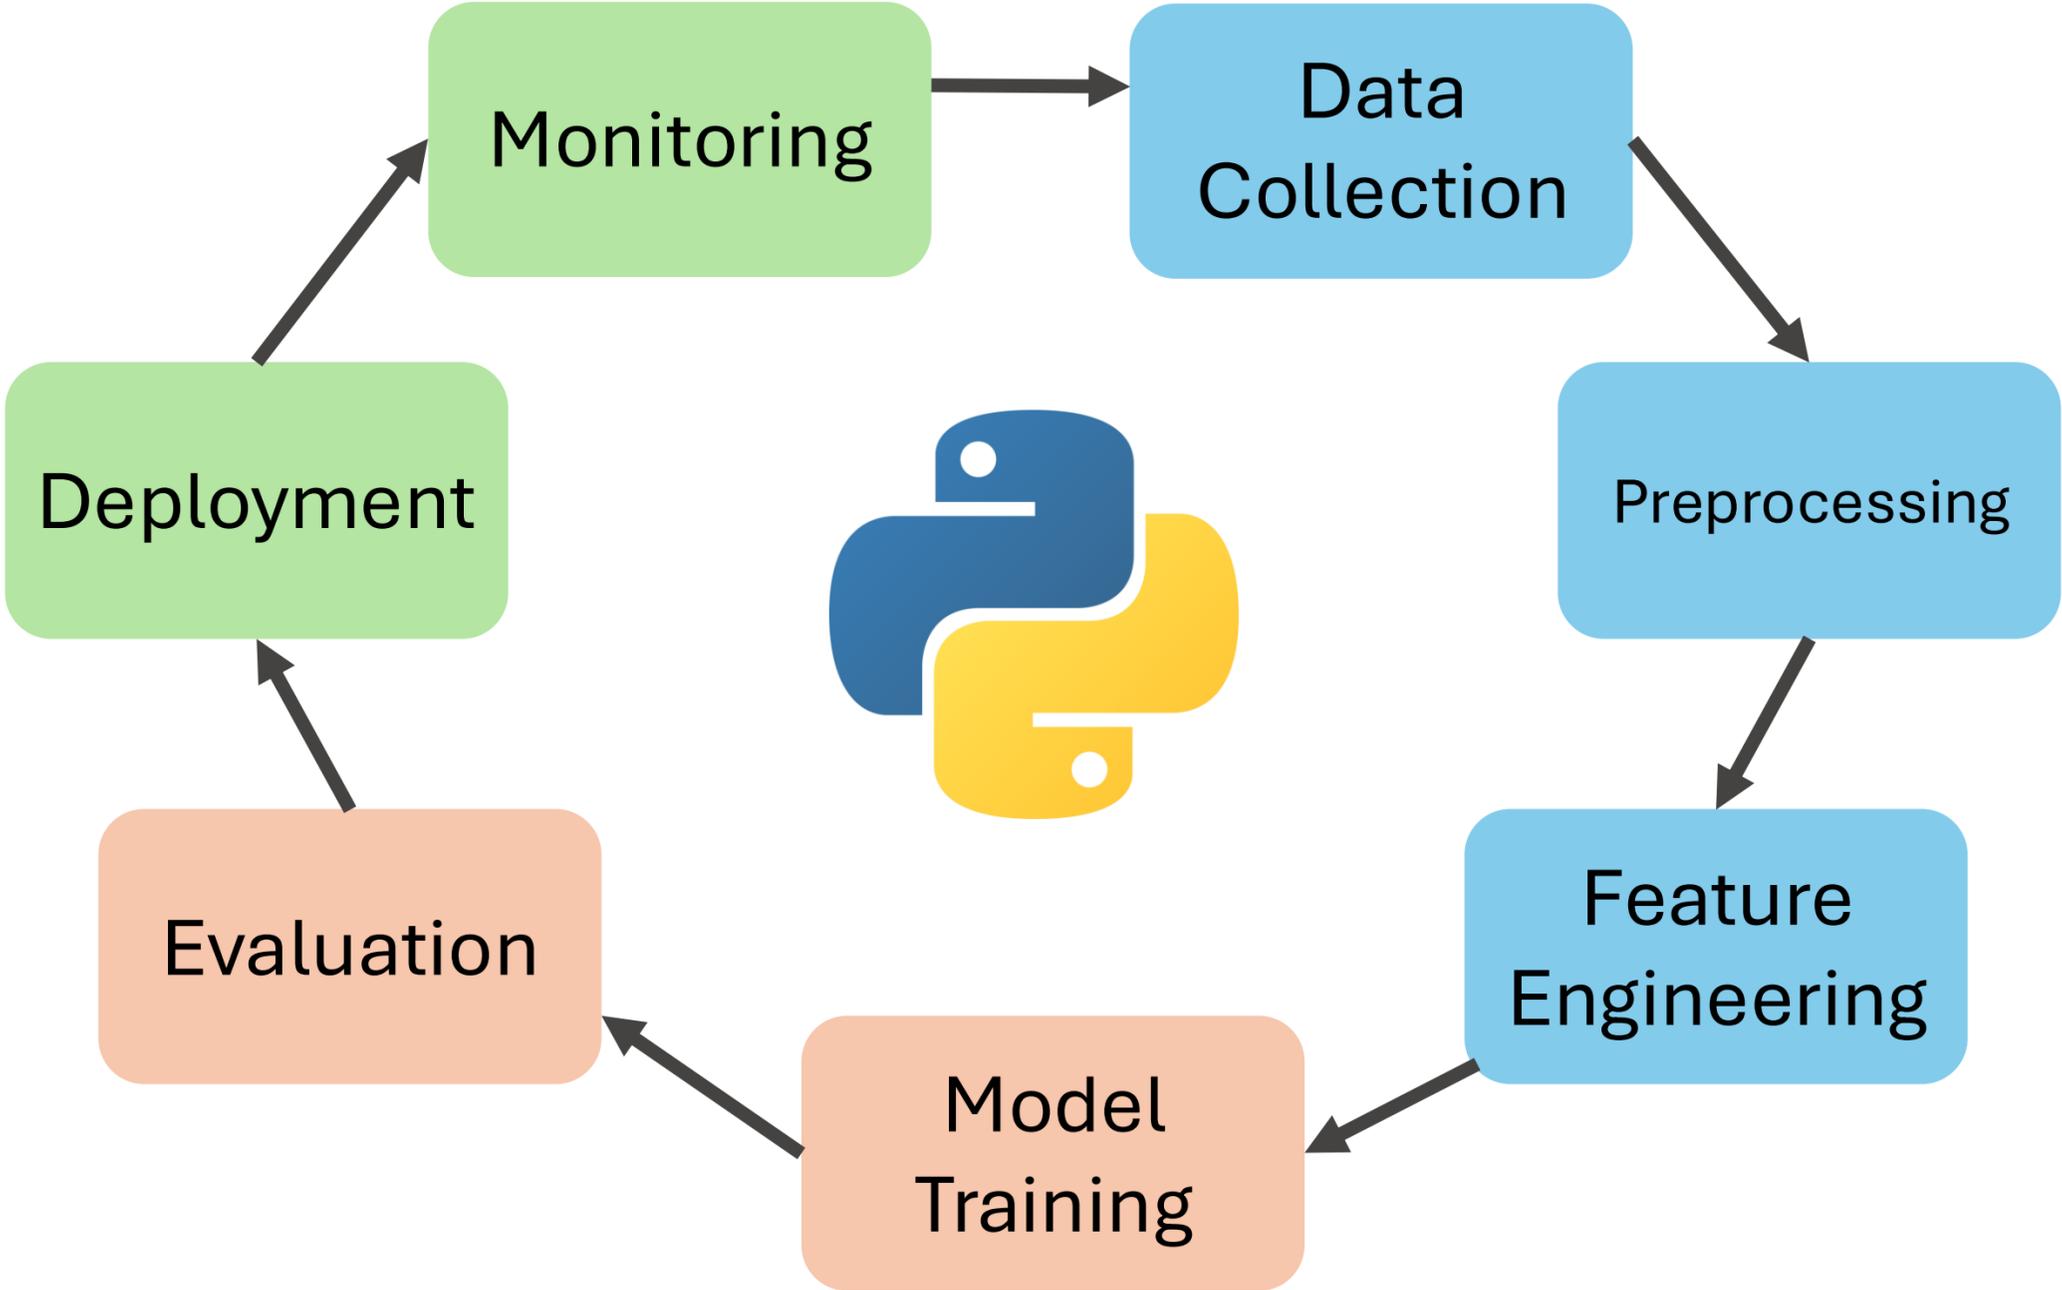
+

=

Machine
Learning

Computational
Techniques

Machine Learning ~~Lifecycle~~ Pipeline



ML Operation

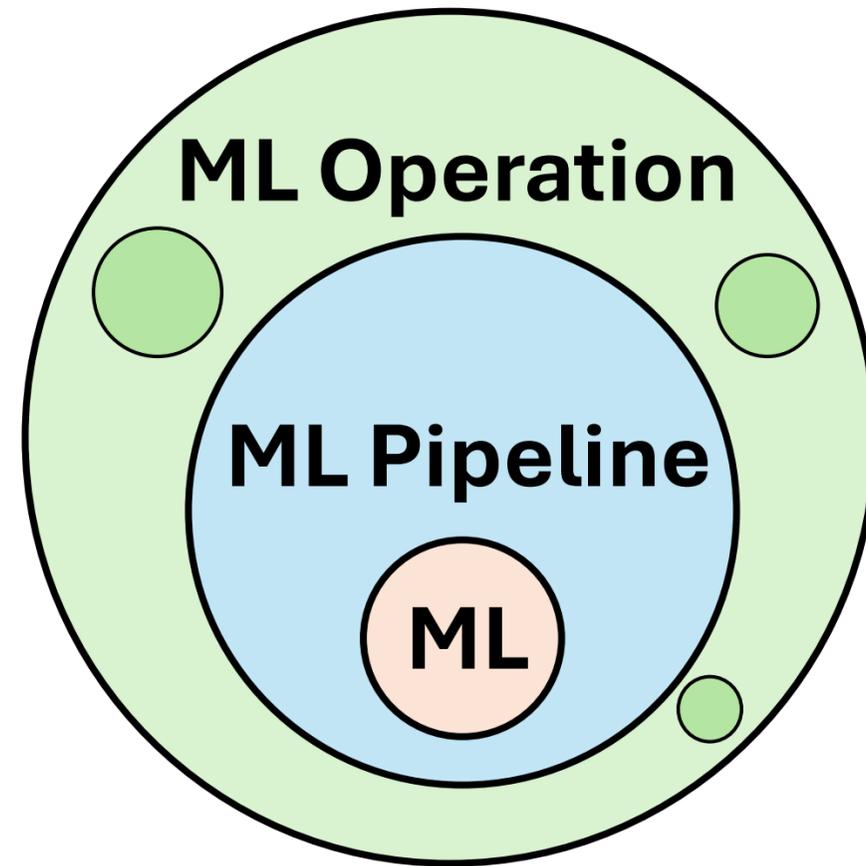
ML Pipeline

ML

Tracking Servers

Compute Clusters

Virtual Environments



Before Fabric



Limitations

- Not a server
- Technical overhead
- Zero redundancy



Microsoft Fabric

- All Tools, One Platform
- Reduced Complexity
- Scalability
- Microsoft Ecosystem

MLOps on Fabric

Basic Framework



x1 Lakehouse



x1 Pipeline



x6 Notebooks



x1 Environment



x2 Experiments



Lakehouse

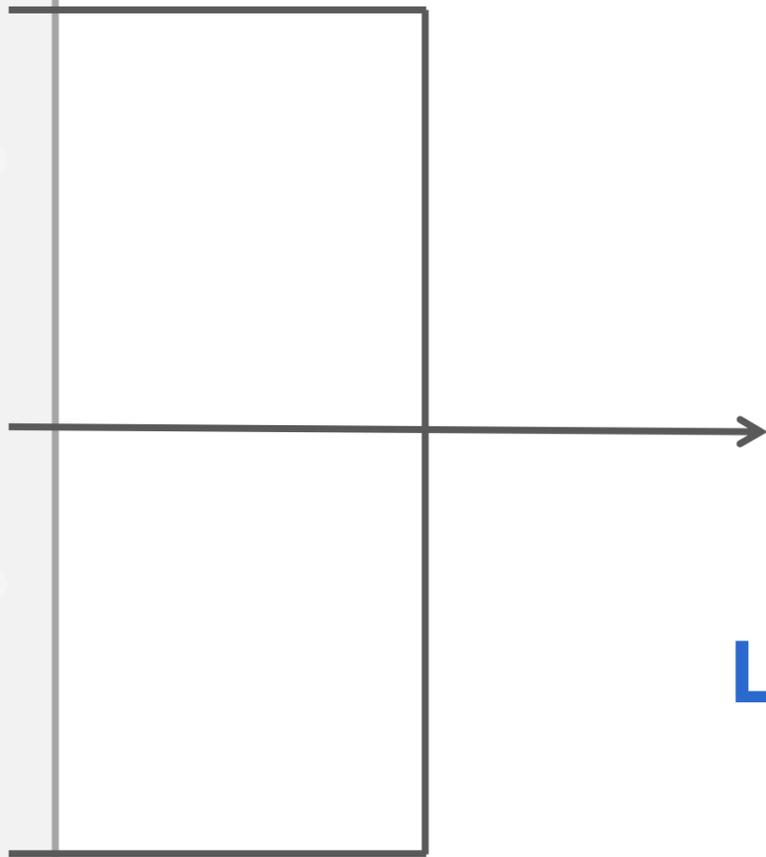
- Separation of Concerns (SoC)

External Databases

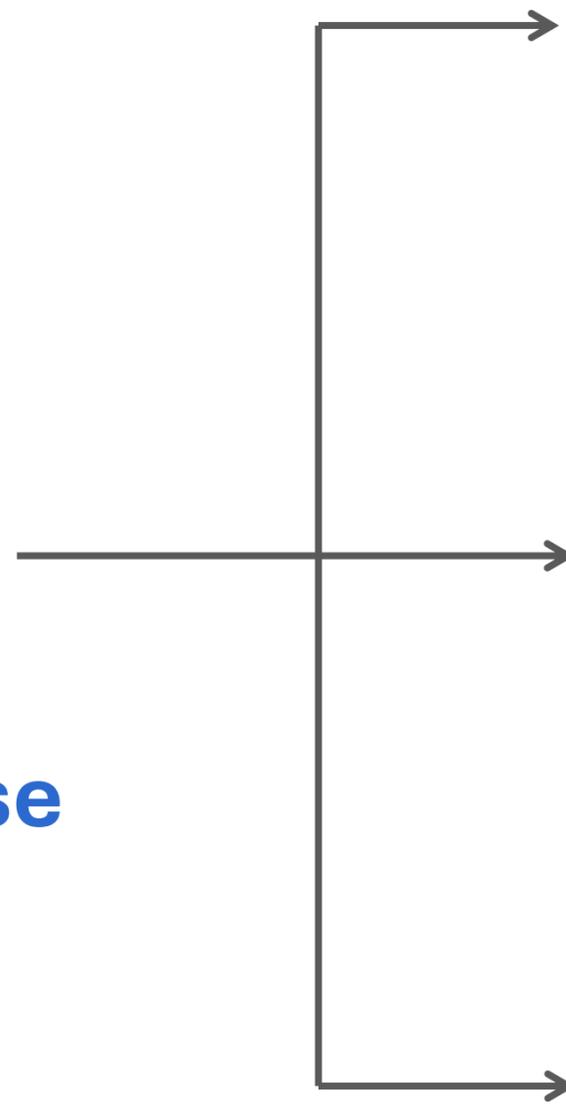


ML Pipeline

External Databases



Lakehouse



EDA



ML Pipeline



Development



Lakehouse

- Separation of Concerns (SoC)
- Reference Tables

11.0101 → Computer Science



CIP Crosswalk

	ABC cip_code	ABC cip_title
614	cip_15_0404	Instrumentation Technology/Technician.
615	cip_15_0405	Robotics Technology/Technician.
616	cip_15_0406	Automation Engineer Technology/Technician.
617	cip_15_0407	Mechatronics, Robotics, and Automation Engineering Te
618	cip_15_0499	Electromechanical Technologies/Technicians, Other.
619	cip_15_05	Environmental Control Technologies/Technicians.
620	cip_15_0501	Heating, Ventilation, Air Conditioning and Refrigeration
621	cip_15_0503	Energy Management and Systems Technology/Technicia
622	cip_15_0505	Solar Energy Technology/Technician.
623	cip_15_0506	Water Quality and Wastewater Treatment Management



Lakehouse

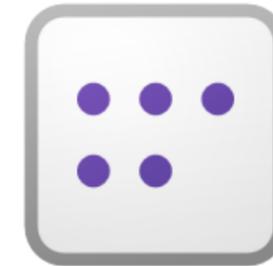
- Separation of Concerns (SoC)
- Reference Tables
- Store Predictions



current_predictions



prediction_history





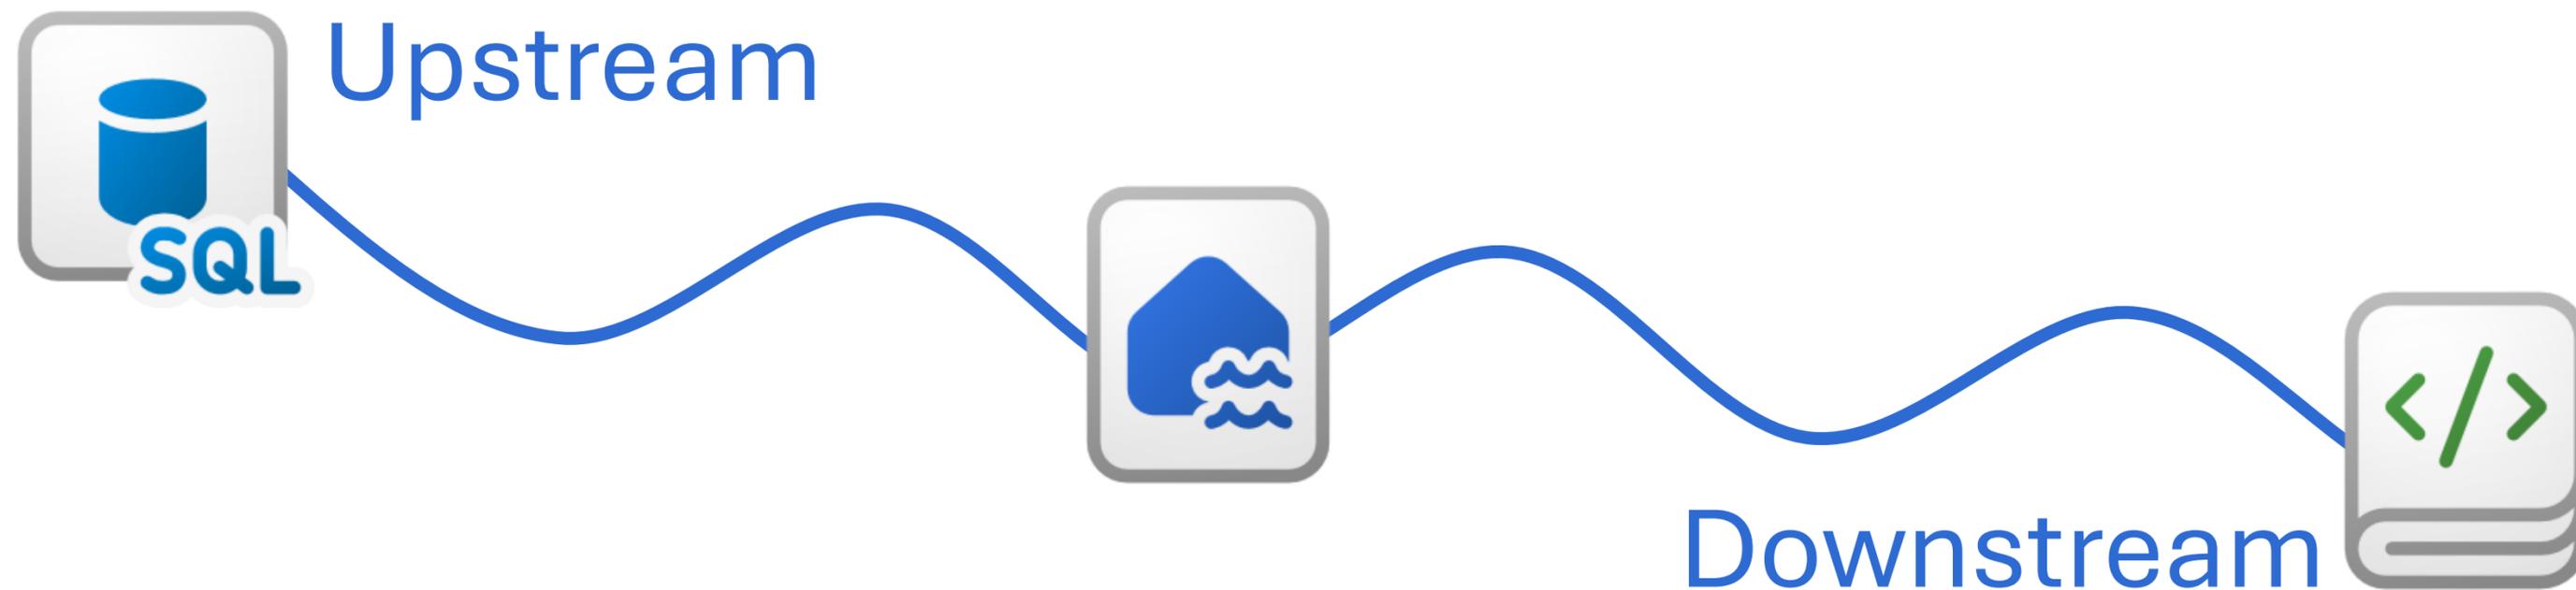
Lakehouse

- Separation of Concerns (SoC)
- Reference Tables
- Store Predictions
- Unstructured Data
 - Configuration files
 - CSV files

Data for ML

- Subset of clean data
- Numerous transformations
 - Feature selection
 - Type casting
 - Encoding
 - Feature scaling
 - Imputation

← Simple Transformations



Complex Transformations →

External Databases

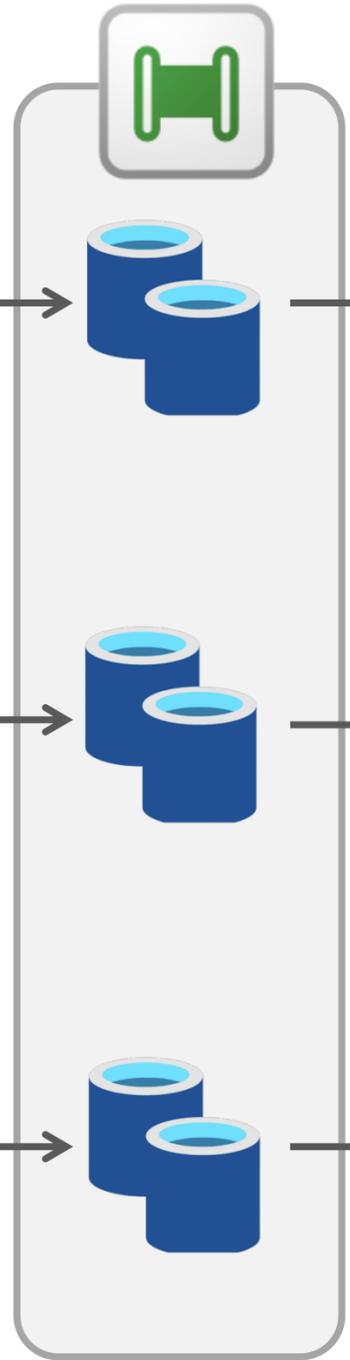


Lakehouse



ML Pipeline

External Databases



 = Copy data activity





Copy data activity

- Extract External Data
- Custom SQL Query
- String Interpolation

```
1 SELECT
2     hs_gpa,
3     ap_credits,
4     act_math_score,
5     cip_code,
6     midterm_gpa,
7     CASE WHEN first_gen = 'Y' THEN 1 ELSE 0 END AS is_first_gen
8 FROM student_enrollment
9 WHERE
10    academic_level = 'freshman'
11    AND semester = 'fall'
12    AND week = 2
13    AND academic_load = 'F'
14
```

```
1  SELECT
2      hs_gpa,
3      ap_credits,
4      act_math_score,
5      cip_code,
6      midterm_gpa,
7      CASE WHEN first_gen = 'Y' THEN 1 ELSE 0 END AS is_first_gen
8  FROM student_enrollment
9  WHERE
10     academic_level = 'freshman'
11     AND semester = 'fall'
12     AND week = '@{variables('week')}'
13     AND academic_load = 'F'
14
```



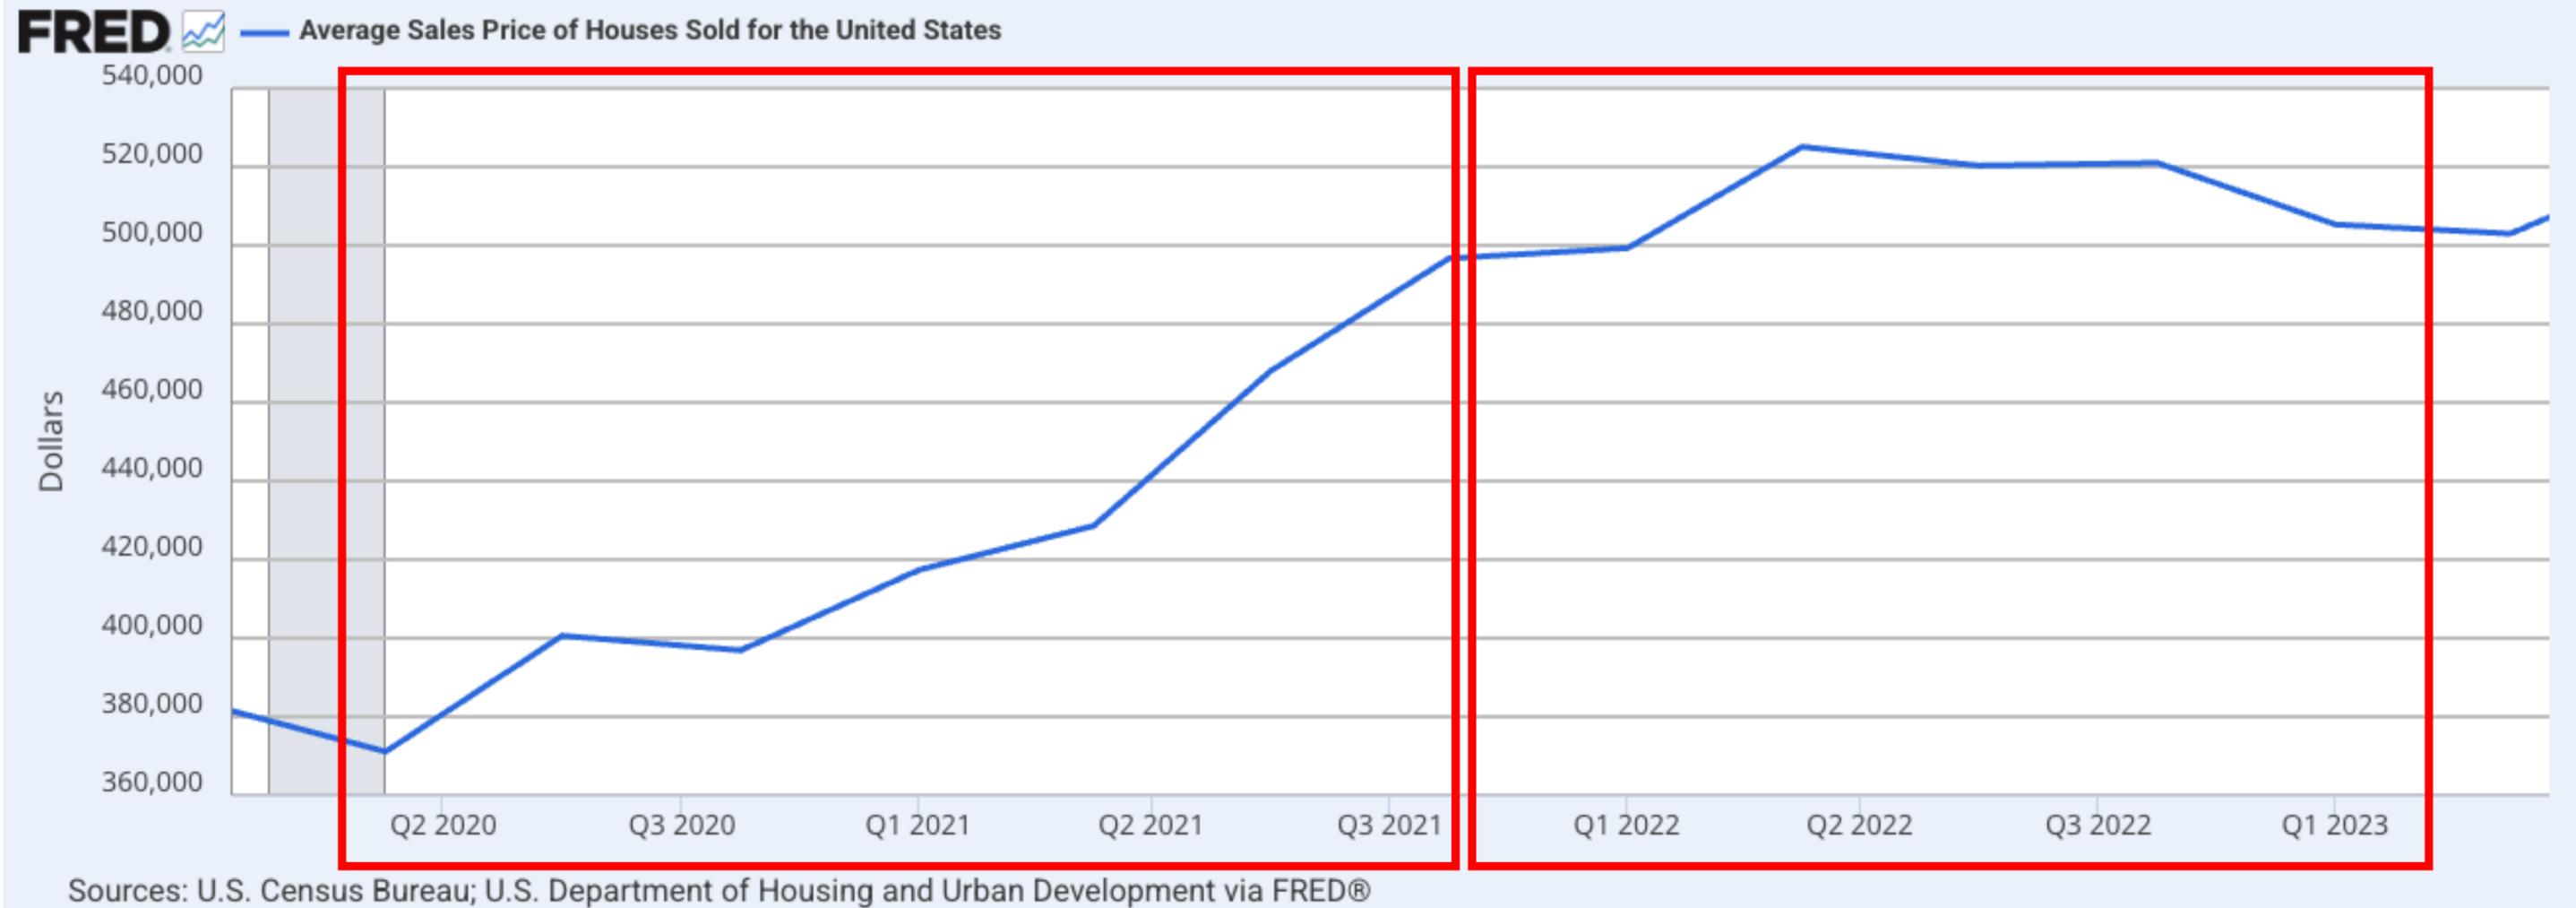
Copy data activity

- Extract External Data
- Custom SQL Query
- String Interpolation
- Explicit Column Mapping



Inferred Schema		Target Schema	
STUDENT_ID	String ✓	STUDENT_ID	String ✓
ACT_MATH	String ✗	ACT_MATH	Int32 ✗





Data Validation

- Many checks per column
- Difficult to maintain

Column Registry

- YAML file
- Column metadata

```
- group name: standardized_test_scores
columns:
  - name: act_math
    display_name: ACT Math Score
    role: feature
    dtype: numeric
    validation:
      range:
        min: 1
        max: 36
      non_null: false
    drift:
      monitor: true
      method: ks
      warn_threshold: 0.1
      alert_threshold: 0.2
    enabled: true
```

← Structural Validation

← Distributional Validation

Basic Framework



x1 Lakehouse



x1 Pipeline



x6 Notebooks



x2 Experiments



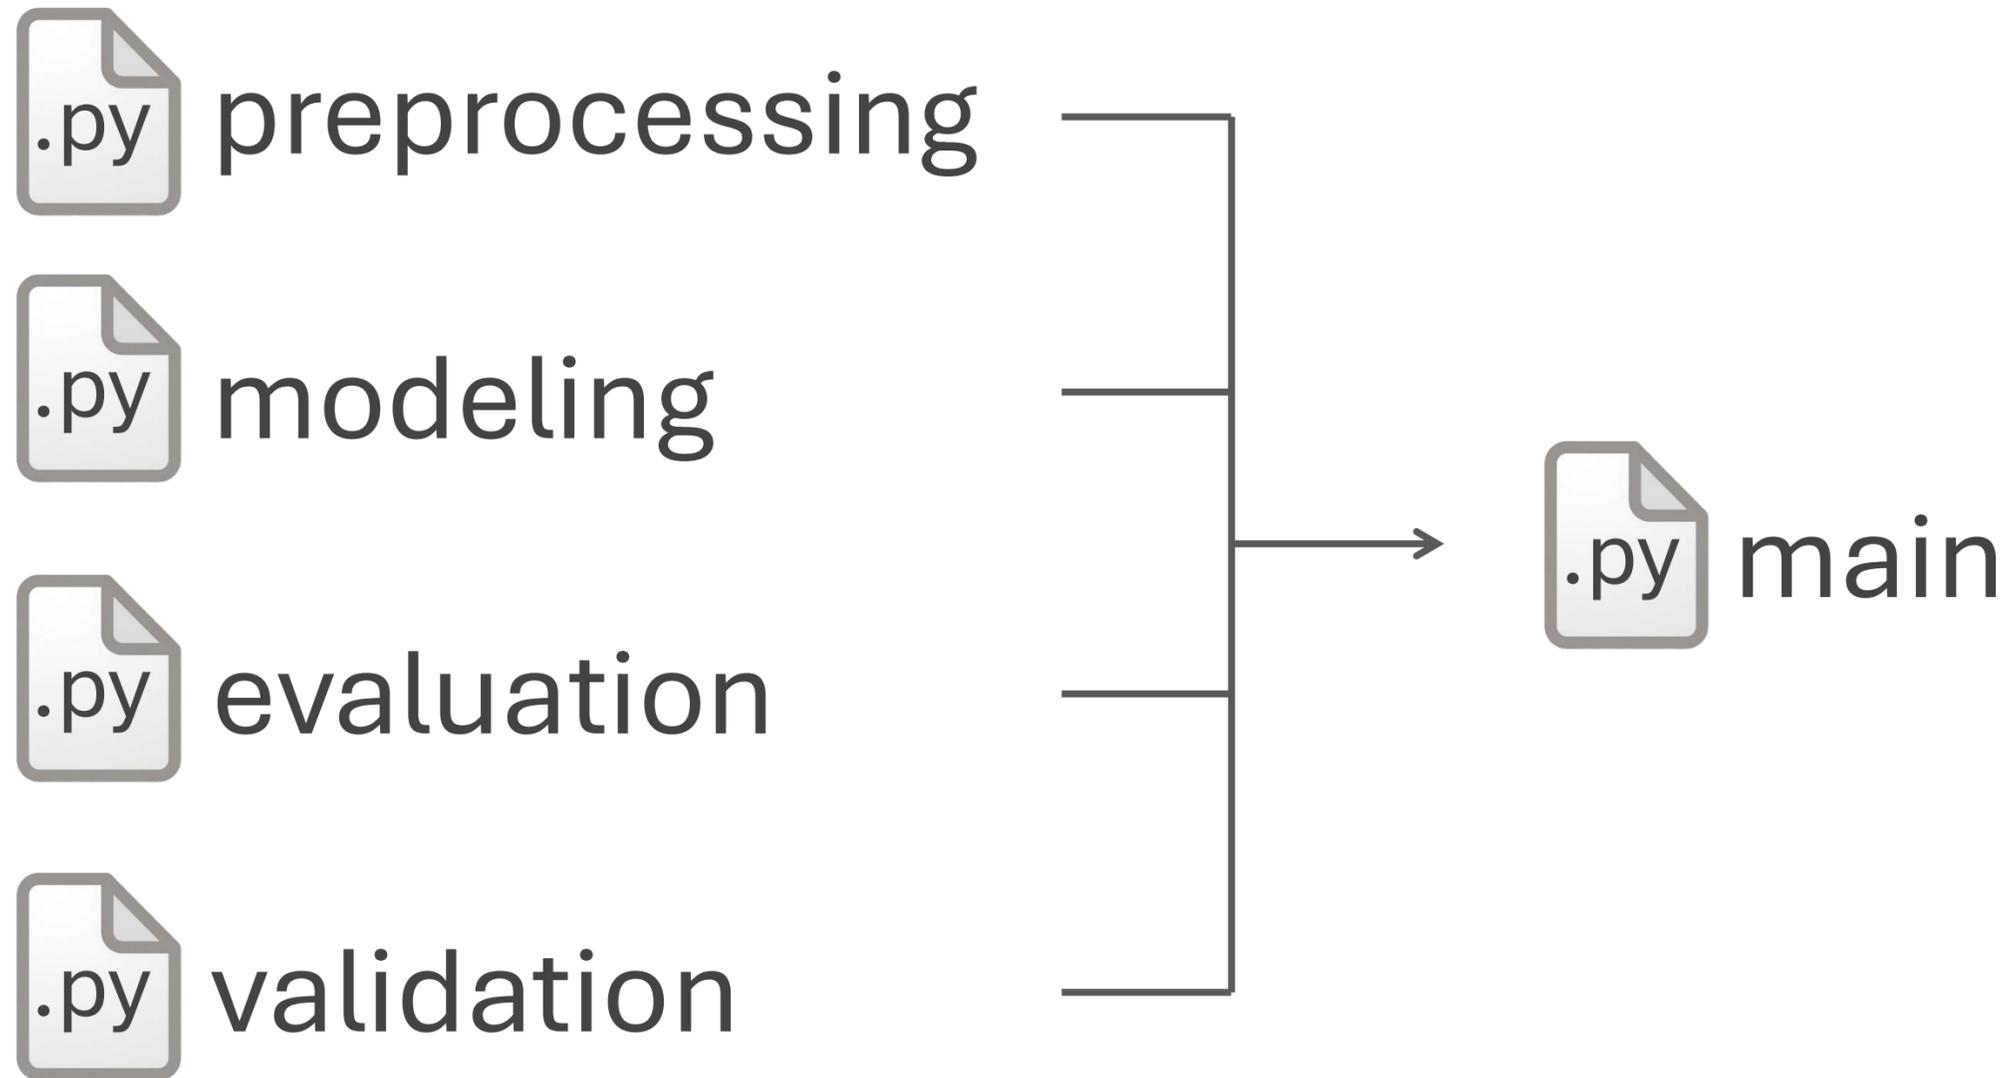
x1 Environment



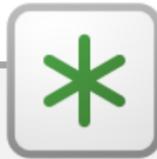
ML Pipeline

- Software
- Python
- Package





Spark Job Definition



-  main
-  preprocessing
-  modeling
-  evaluation
-  validation

Notebook



-  main
-  preprocessing
-  modeling
-  evaluation
-  validation

Hybrid



-  main
-  preprocessing
-  modeling
-  evaluation
-  validation



Main definition file * ⓘ

Upload or enter the url of a .py file. File size must be less than 250MB.

- Upload a local file
 ADLS-Gen2 path

e.g. "abfss://.../path/to/shakespeare.txt abfss://.../path/to/result"



Reference file ⓘ

Upload or enter the urls of one or multiple files. File size must be less than 250MB.

- Upload local files
 ADLS-Gen2 path

Each line should be a storage URI. e.g.
"abfss://.../path/to/shakespeare.txt abfss://.../path/to/result"



Command line arguments ⓘ

e.g. "abfss://.../path/to/shakespeare.txt abfss://.../path/to/resi"

Lakehouse Reference * ⓘ

+ Add

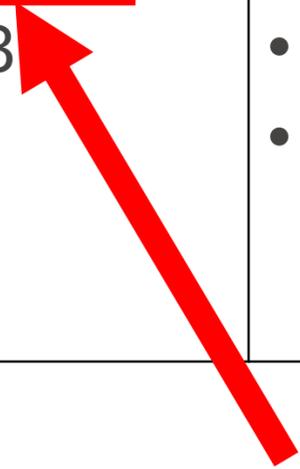


∨ All lakehouses (0)



Environment (Default)

Runtime	Libraries	Compute
<ul style="list-style-type: none">• OS Mariner 2.0• Apache Spark 3.5• Delta Lake 3.2• Python 3.11	<ul style="list-style-type: none">• scikit-learn 1.2.2• xgboost 2.0.3• shap 0.42.1	<ul style="list-style-type: none">• Node size: medium• # of Nodes: 1-10• Autoscale: enabled



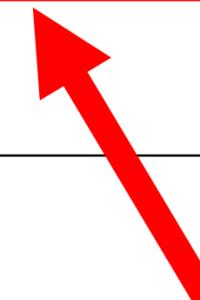
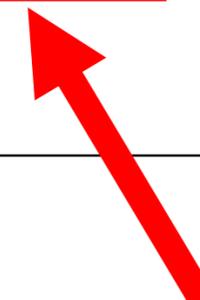
Machine Learning Libraries

- Single-node libraries
 - Scikit-Learn
 - XGBoost
- Multi-node libraries
 - Synapse ML
 - MLlib

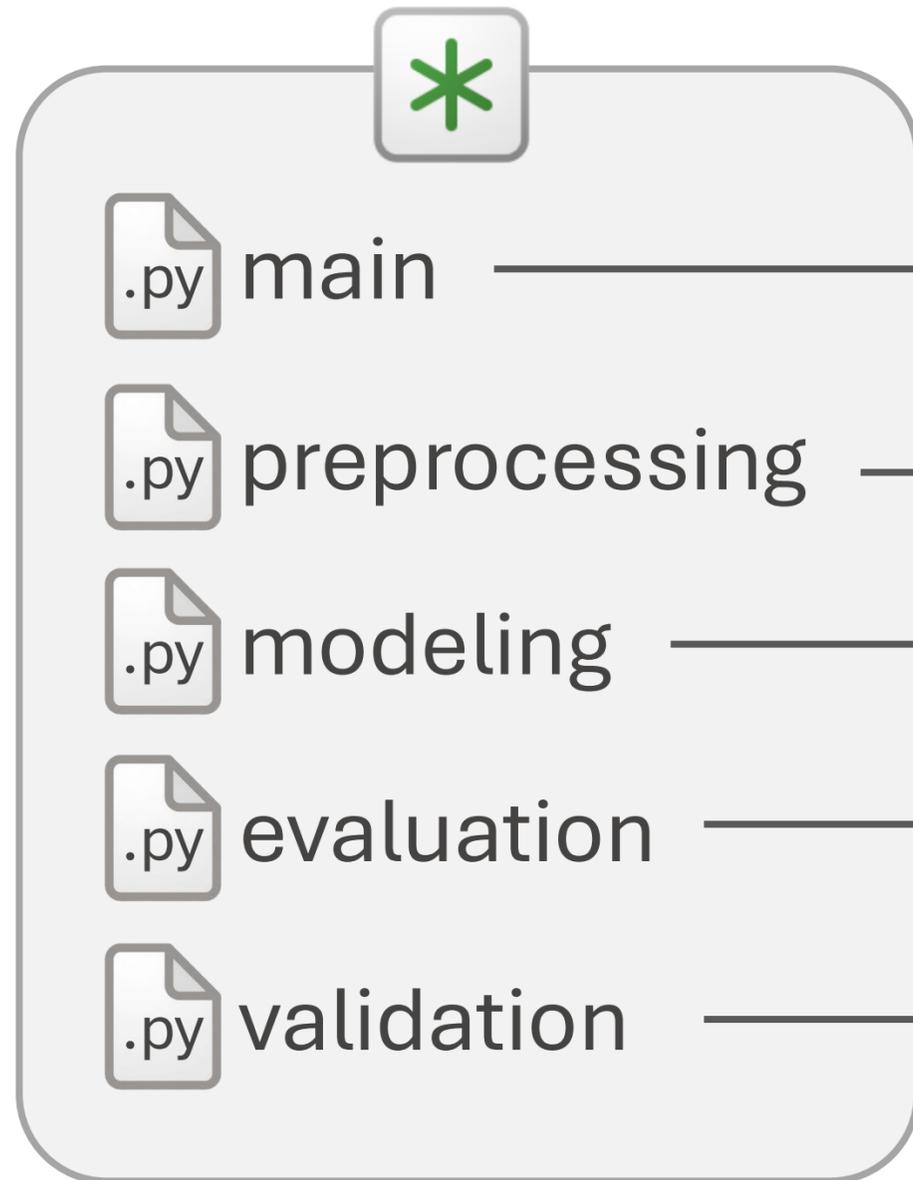


Environment (Custom)

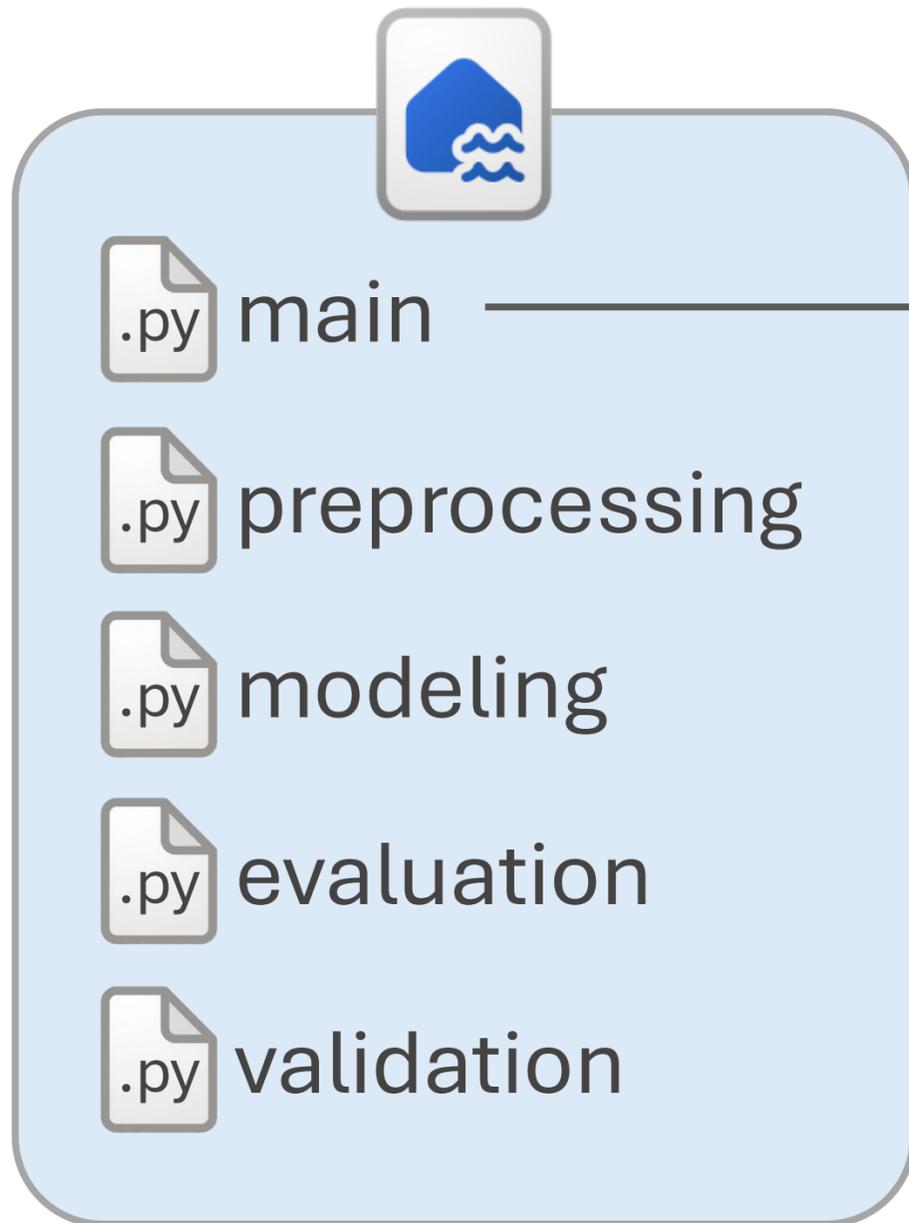
Runtime	Libraries	Compute
<ul style="list-style-type: none">• OS Mariner 2.0• Apache Spark 3.5• Delta Lake 3.2• Python 3.11	<ul style="list-style-type: none">• scikit-learn 1.6.0• xgboost 3.0.0• shap 0.48.0	<ul style="list-style-type: none">• Node size: large• # of Nodes: 1• Autoscale: disabled



Spark Job Definition



 Environment (Default)		
Runtime	Libraries	Compute
<ul style="list-style-type: none">• OS Mariner 2.0• Apache Spark 3.5• Delta Lake 3.2• Python 3.11	<ul style="list-style-type: none">• scikit-learn 1.2.2• xgboost 2.0.3• shap 0.42.1	<ul style="list-style-type: none">• Node size: medium• # of Nodes: 1-10• Autoscale: enabled



```
1 import sys
2 import os
3
4 # Add directory to your Python path
5 sys.path.append("/lakehouse/default/Files/")
```

PySpark (Python) ▾

```
1 %load file:/lakehouse/default/Files/main.py
```

PySpark (Python) ▾



main



preprocessing



modeling



evaluation

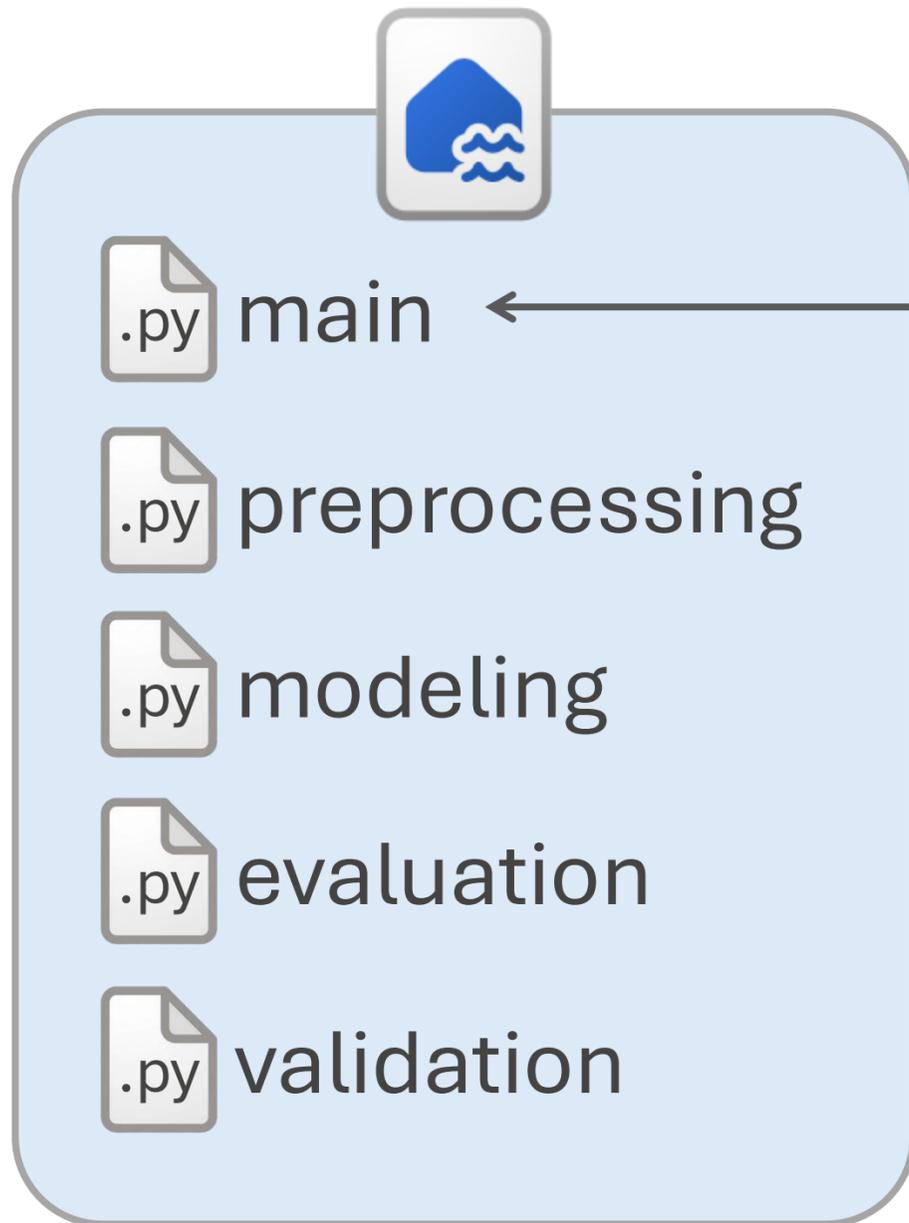


validation

```
1 import sys
2 import os
3
4 # Add directory to your Python path
5 sys.path.append("/lakehouse/default/Files/")
```

PySpark (Python) ▾

```
1 # %load file:/lakehouse/default/Files/main.py
2 import modeling
3 def main():
4     spark = SparkSession.builder.getOrCreate()
```



```
1 import sys
2 import os
3
4 # Add directory to your Python path
5 sys.path.append("/lakehouse/default/Files/")
```

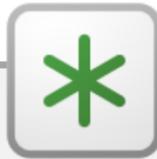
PySpark (Python) ▾

```
1 %%writefile file:/lakehouse/default/Files/main.py
2 import modeling
3 def main():
4     spark = SparkSession.builder.getOrCreate()
5     print('Hello, Fabcon!')
```



This will create a .crc file in your lakehouse

Spark Job Definition



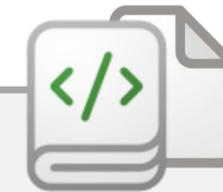
-  main
-  preprocessing
-  modeling
-  evaluation
-  validation

Notebook



-  main
-  preprocessing
-  modeling
-  evaluation
-  validation

Hybrid



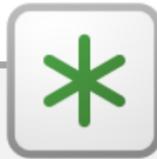
-  main
-  preprocessing
-  modeling
-  evaluation
-  validation



Notebook Approach

- Notebooks in production
- “Primary coding item”
- Convenient
- Use `%run` to import modules into main

Spark Job Definition



-  main
-  preprocessing
-  modeling
-  evaluation
-  validation

Notebook



-  main
-  preprocessing
-  modeling
-  evaluation
-  validation

Hybrid



-  main
-  preprocessing
-  modeling
-  evaluation
-  validation

Hybrid Approach

- Best of both worlds
- Leverage agentic coding tools
 - Claude Code
 - GitHub Copilot
- Perform better on .py files

Experiment Tracking

- Train/test multiple different models
- Analyze each “run”
- Compare different runs

XGB Run

```
with mlflow.start_run():  
    model.fit()  
    mlflow.log_metrics()  
    mlflow.log_params()  
    mlflow.set_tags()  
    mlflow.log_model()
```

Experiment



Run name	Start time	Duration	Accuracy
RF Run	3/19/26	3:13	0.76
XGB Run	3/18/26	4:01	0.78

Experiment



XGB Run

```
with mlflow.start_run():  
    model.fit()  
    mlflow.log_metrics()  
    mlflow.log_params()  
    mlflow.set_tags()  
    mlflow.log_model()
```

RF Run

```
with mlflow.start_run():  
    model.fit()  
    mlflow.log_metrics()  
    mlflow.log_params()  
    mlflow.set_tags()  
    mlflow.log_model()
```

Run name	Start time	Duration	Accuracy
XGB Run	3/20/26	4:25	0.81
RF Run	3/19/26	3:13	0.76
XGB Run	3/18/26	4:01	0.78

Experiment



XGB Run

```
with mlflow.start_run():  
    model.fit()  
    mlflow.log_metrics()  
    mlflow.log_params()  
    mlflow.set_tags()  
    mlflow.log_model()
```

RF Run

```
with mlflow.start_run():  
    model.fit()  
    mlflow.log_metrics()  
    mlflow.log_params()  
    mlflow.set_tags()  
    mlflow.log_model()
```

Run name	Start time	Duration	Accuracy
RF Run	3/20/26	3:11	0.79
XGB Run	3/20/26	4:25	0.81
RF Run	3/19/26	3:13	0.76
XGB Run	3/18/26	4:01	0.78

XGB Run

```
with mlflow.start_run():  
    model.fit()  
    mlflow.log_metrics()  
    mlflow.log_params()  
    mlflow.set_tags()  
    mlflow.log_model()
```

RF Run

```
with mlflow.start_run():  
    model.fit()  
    mlflow.log_metrics()  
    mlflow.log_params()  
    mlflow.set_tags()  
    mlflow.log_model()
```

Experiment



Run name	Start time	Duration	Accuracy
RF Run	3/20/26	3:11	0.79
XGB Run	3/20/26	4:25	0.81
RF Run	3/19/26	3:13	0.76
XGB Run	3/18/26	4:01	0.78

Model



Version 2

Version 1



Basic Framework



x1 Lakehouse



x1 Pipeline



x6 Notebooks



x1 Environment



x2 Experiments

External Database




Custom Environment

-  registry
-  preprocessing
-  modeling
-  evaluation
-  validation

External Database



enrollment



Custom Environment



registry



preprocessing



modeling

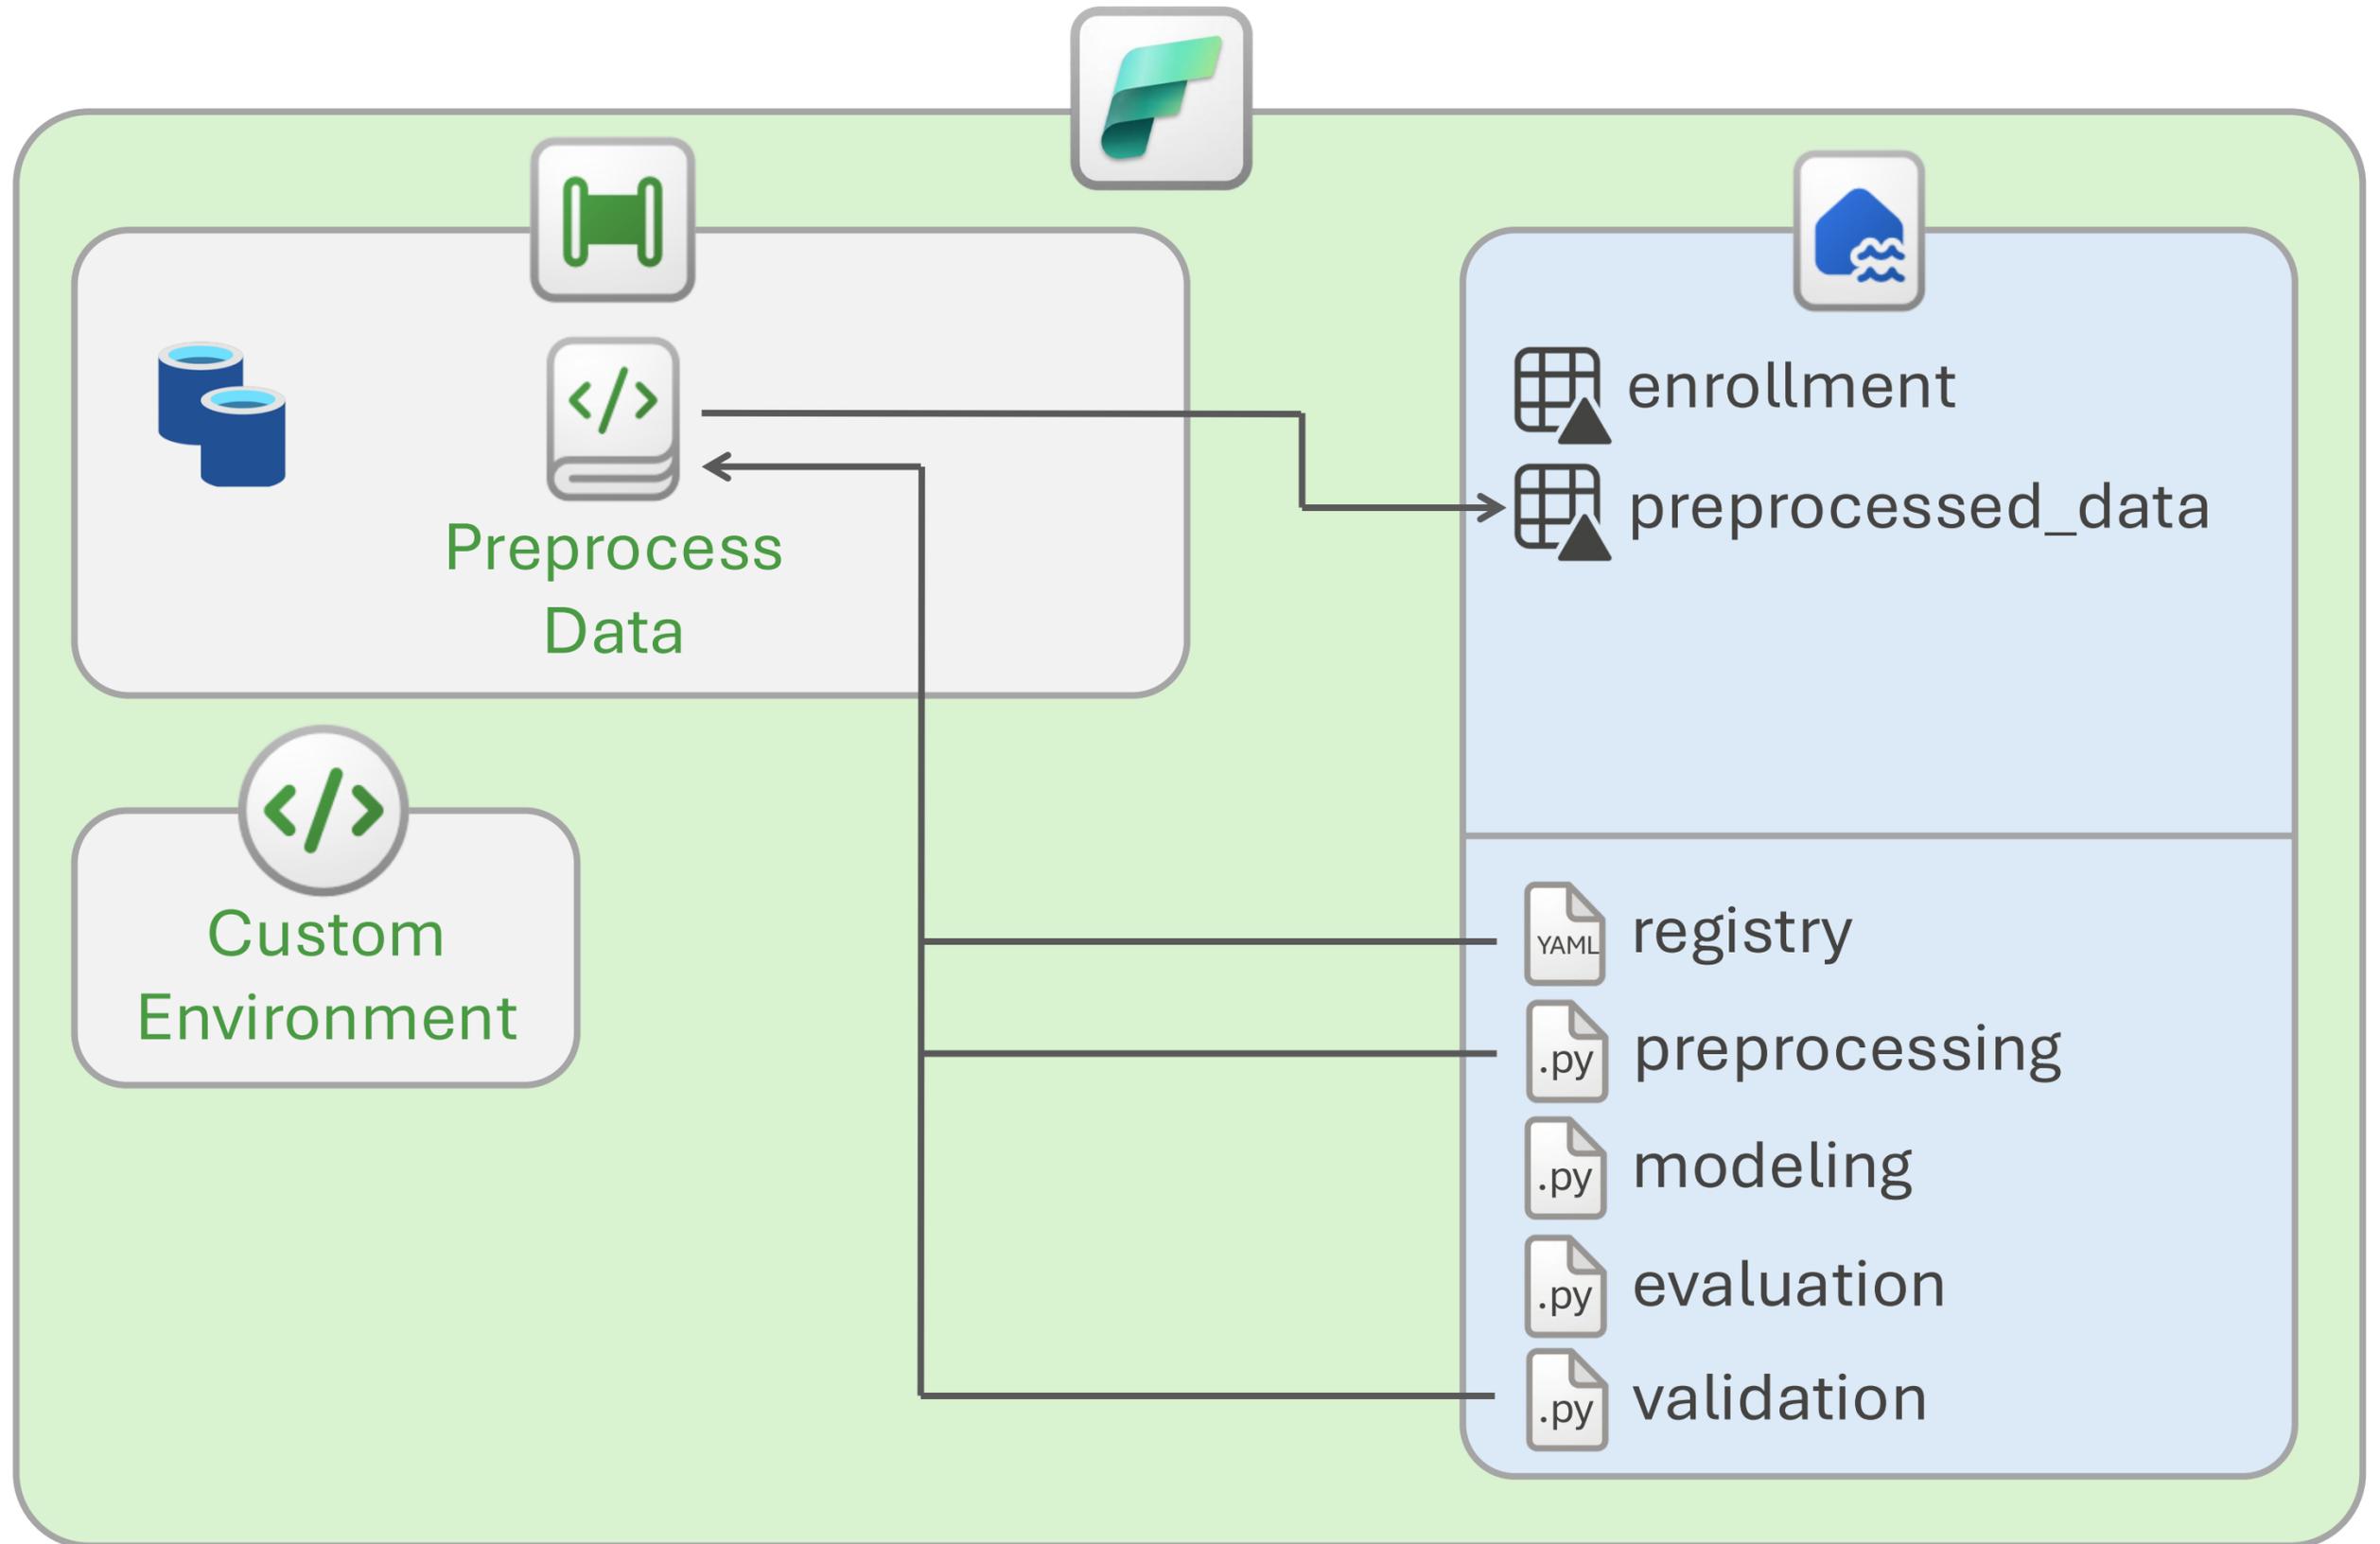


evaluation

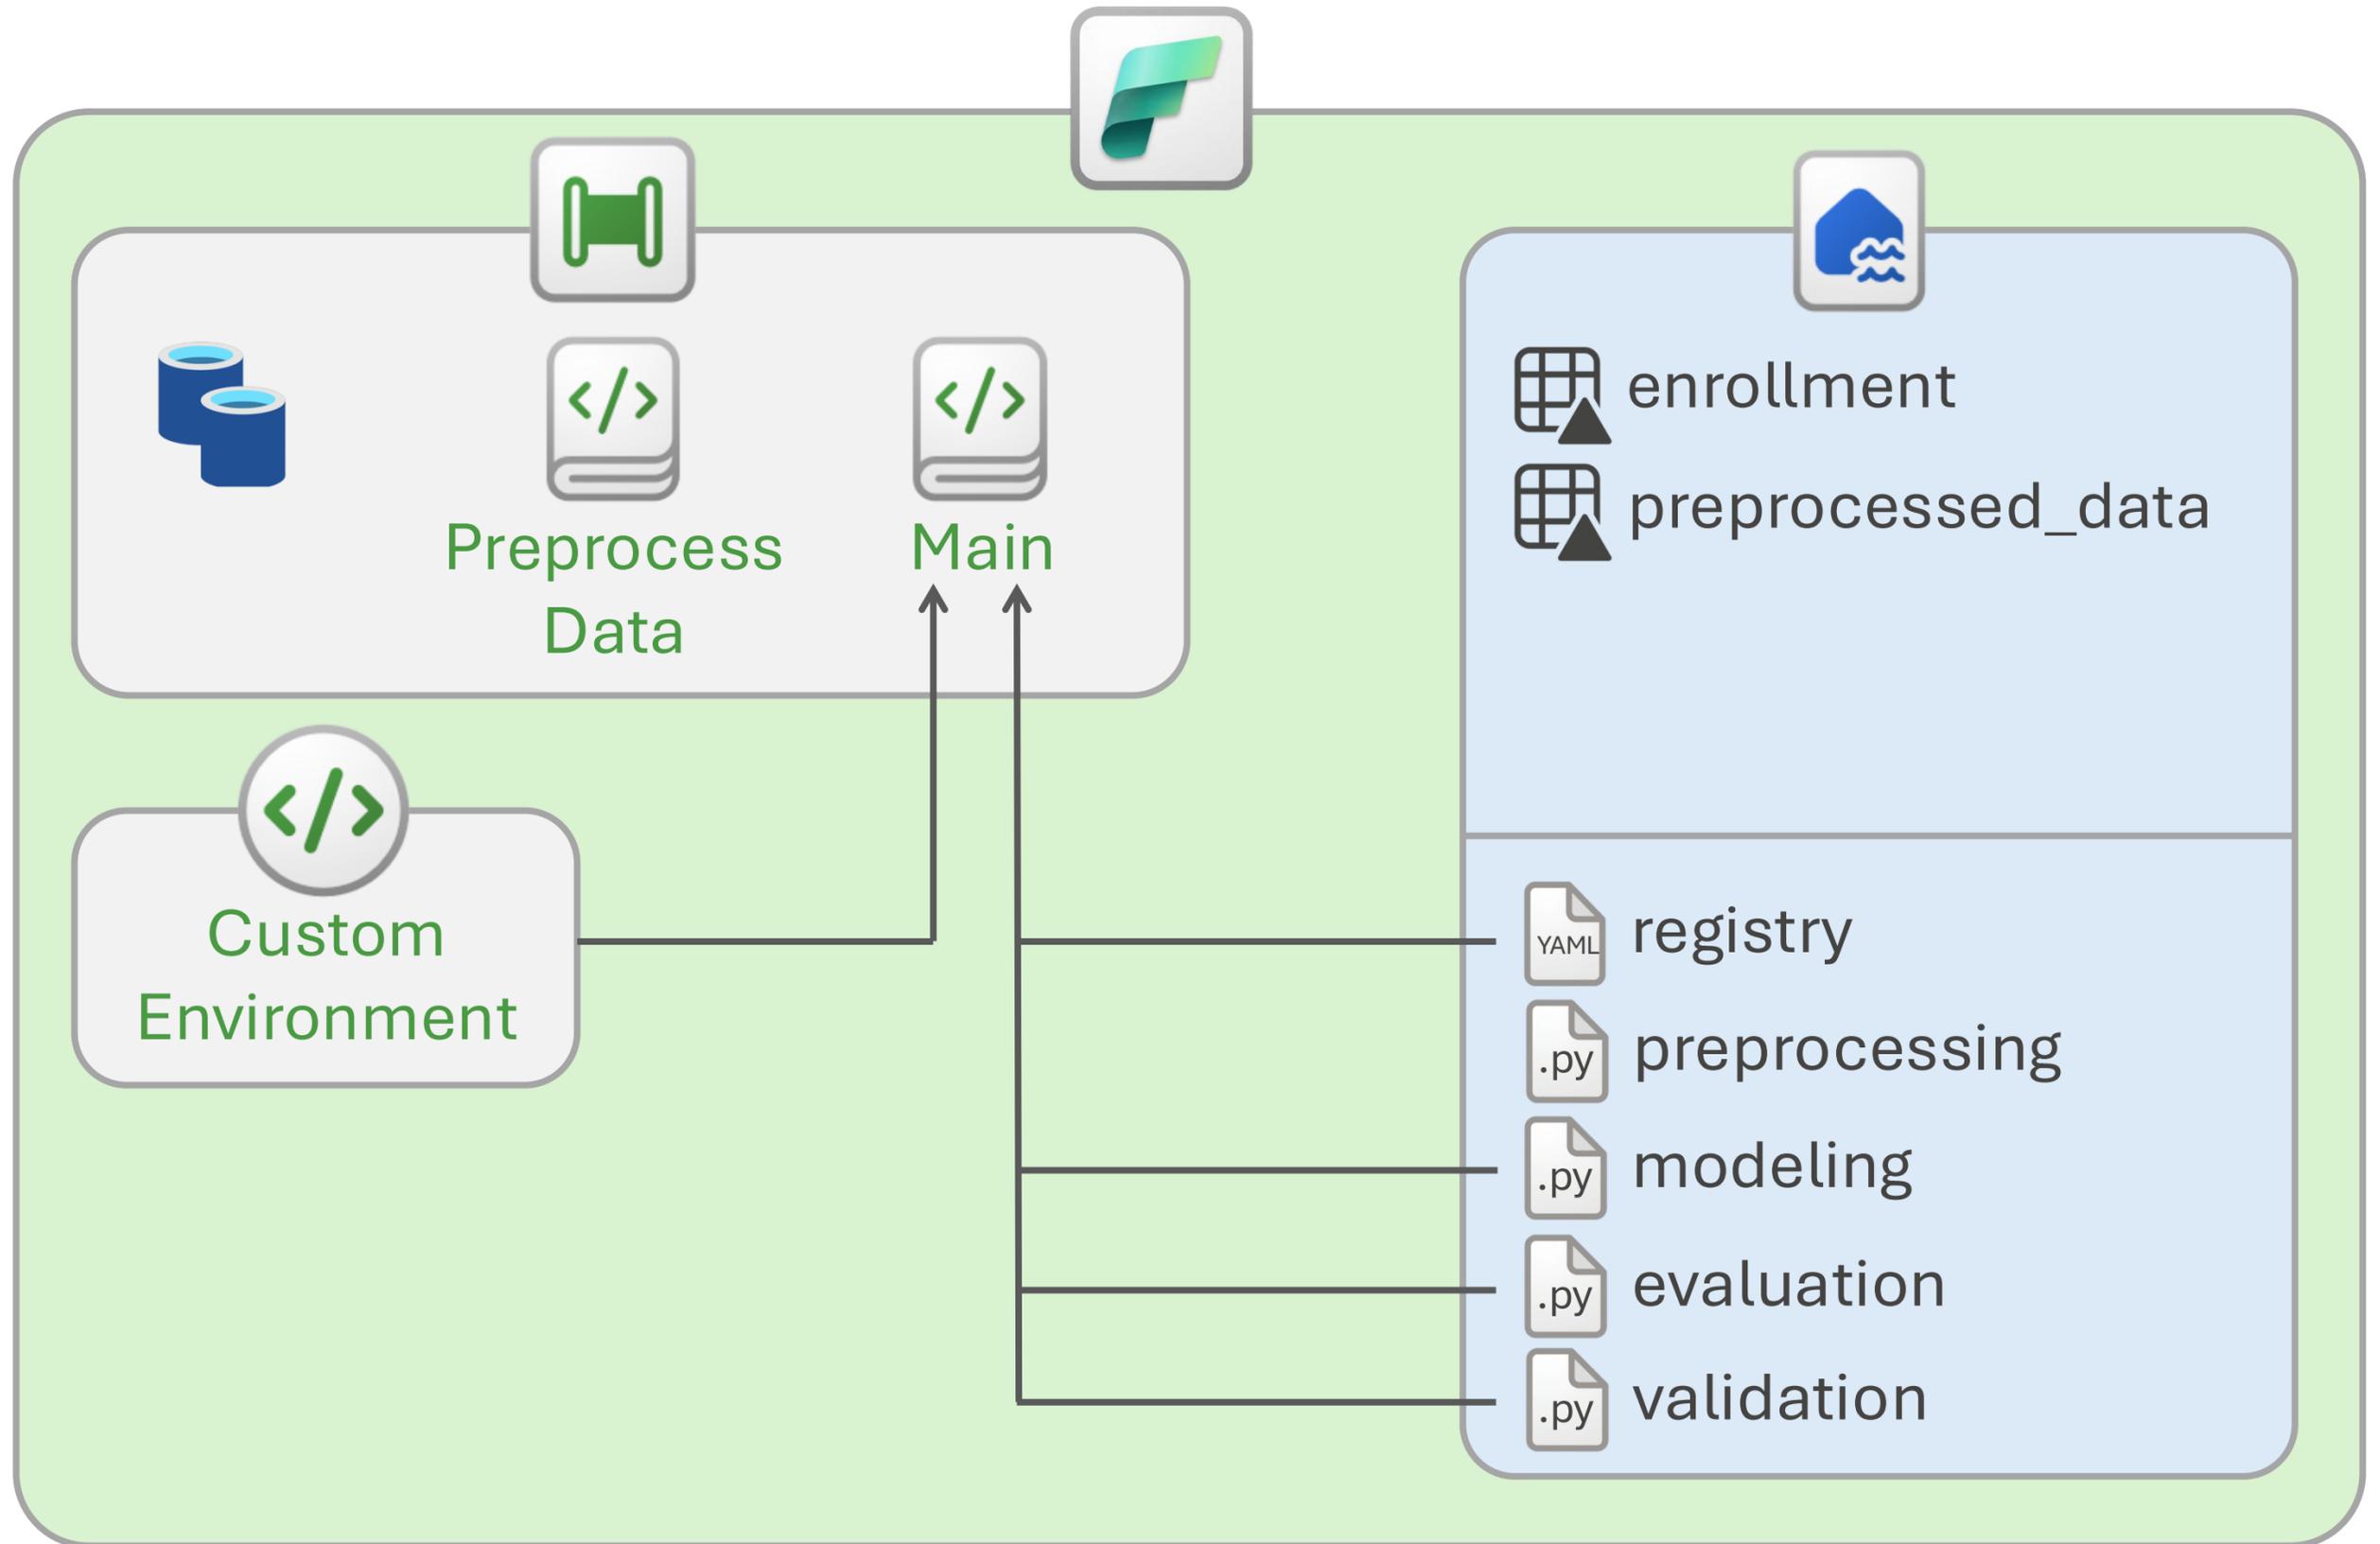


validation

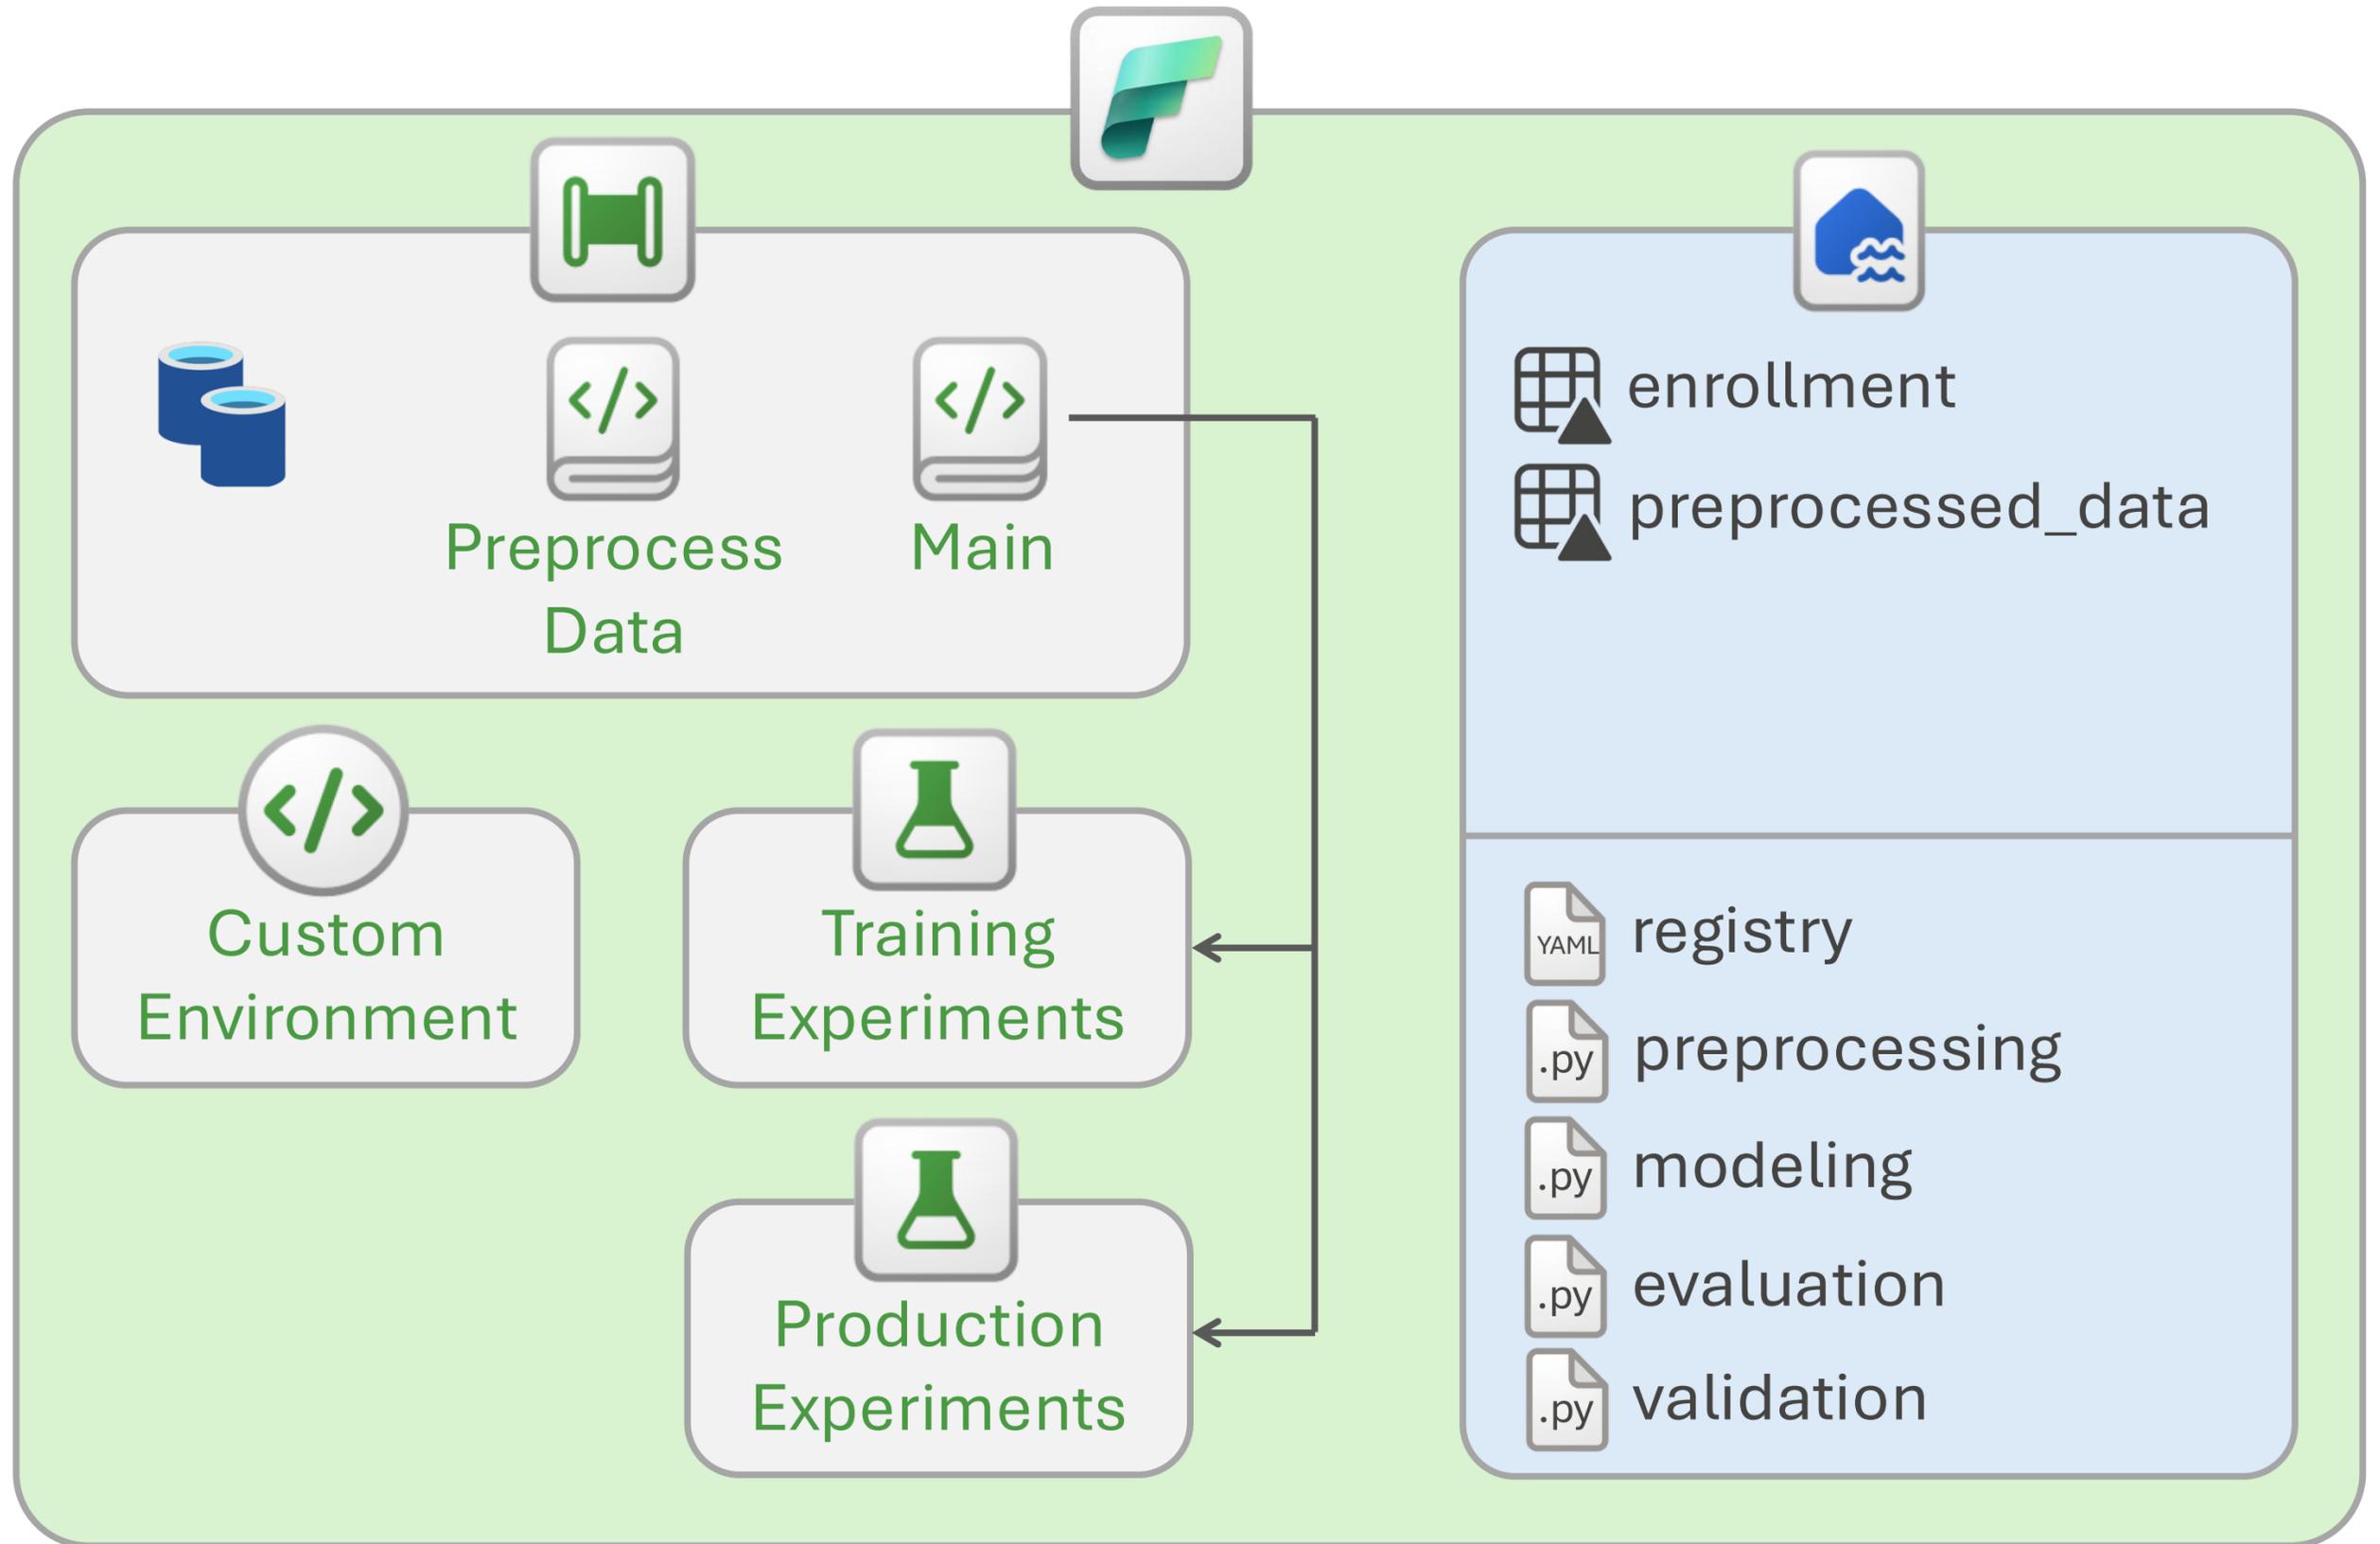
External Database



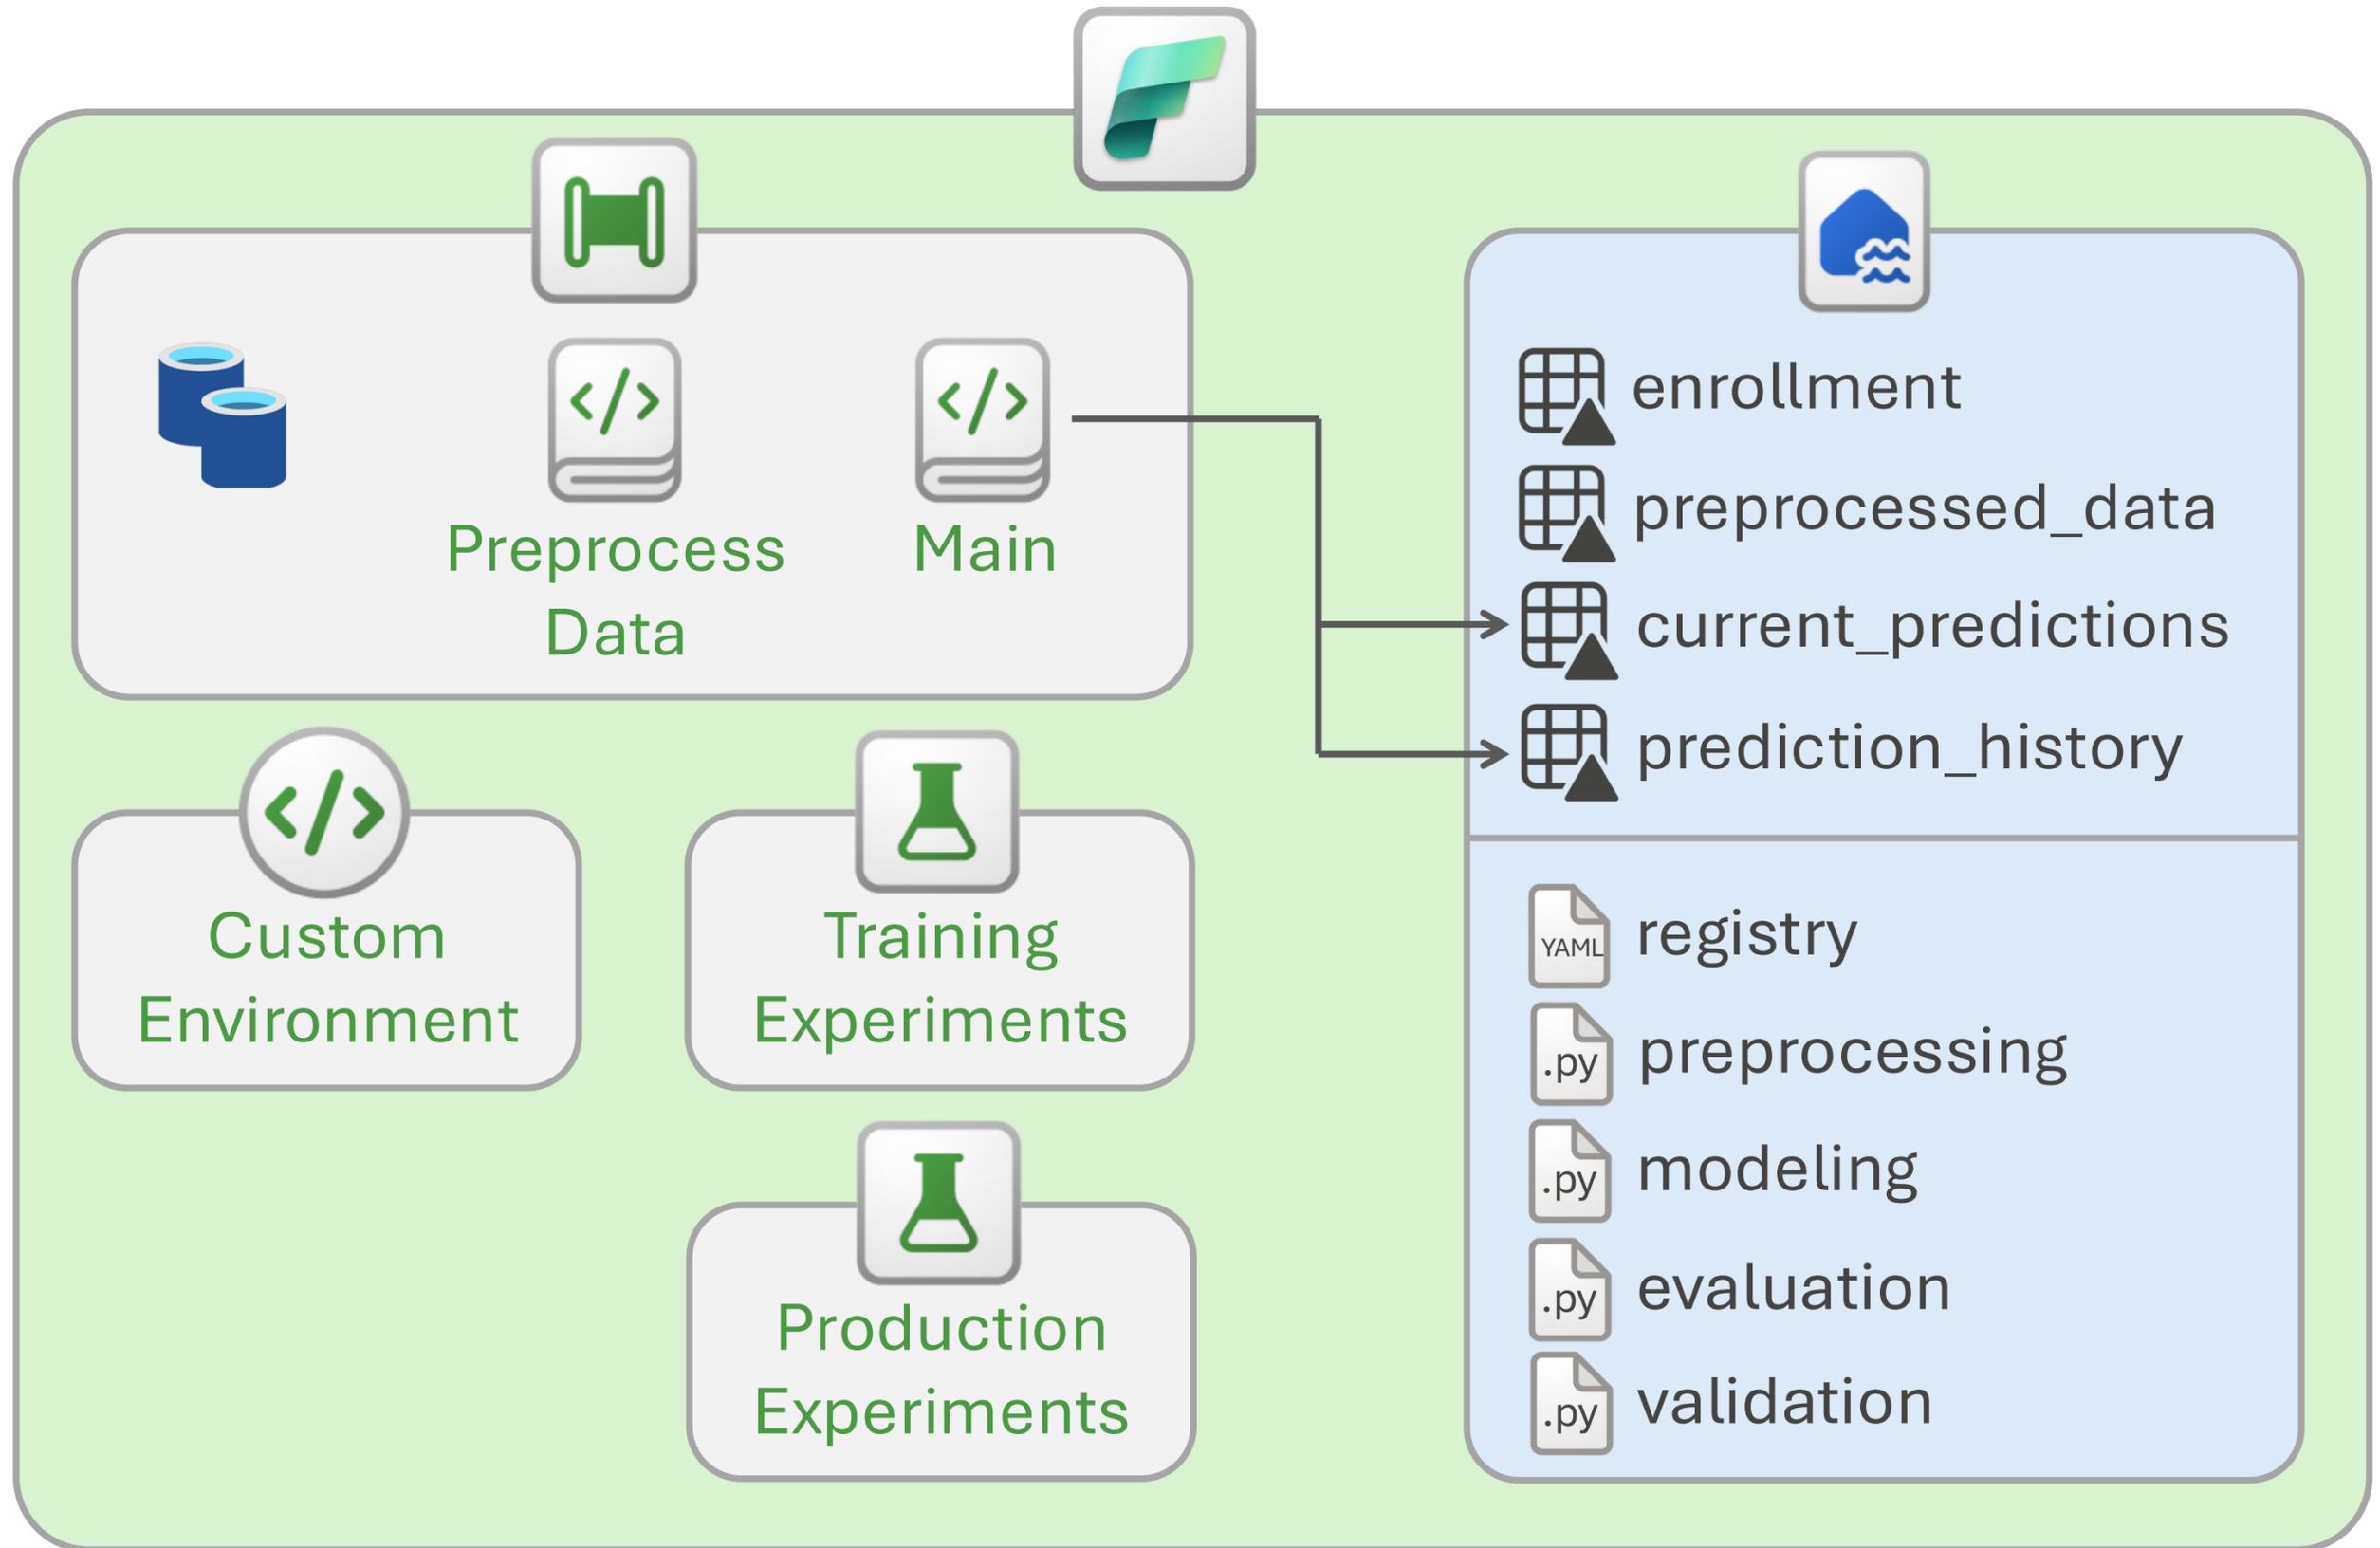
External Database



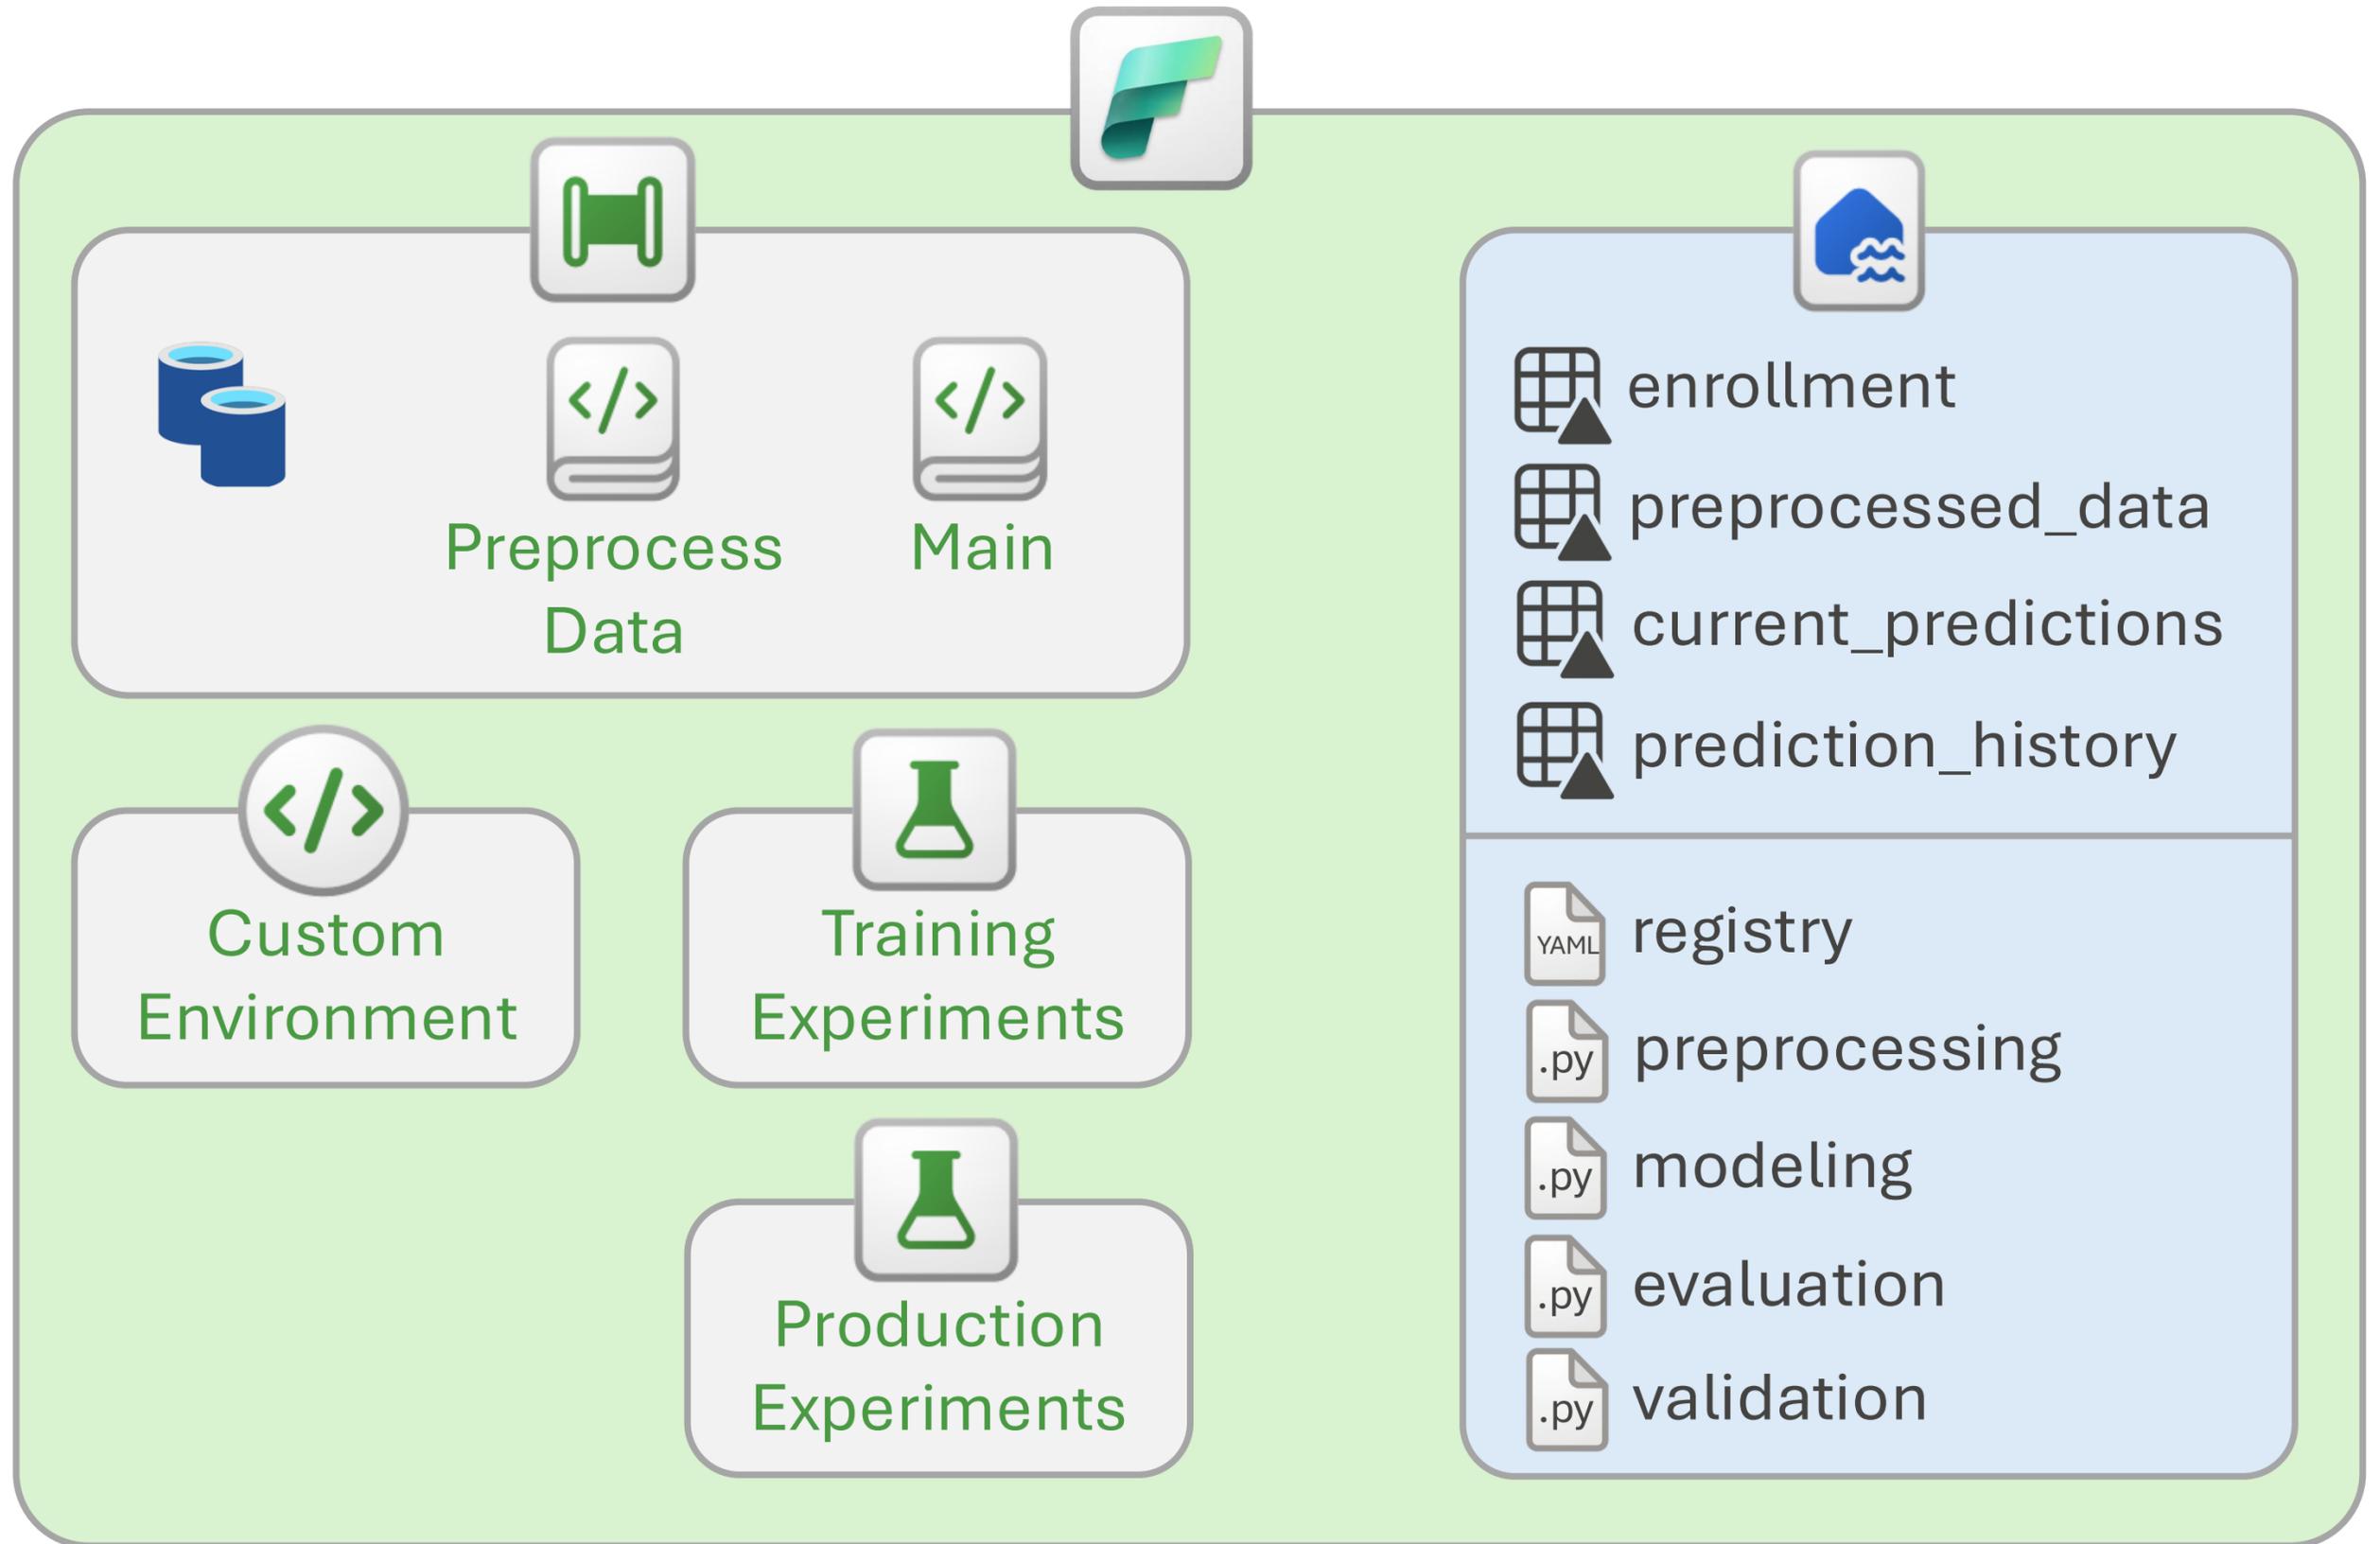
External Database



External Database



External Database



How was the session?



Complete Session Surveys in
Whova for your chance to WIN
PRIZES!

