Enterprise-Scale Machine Learning with IBM Db2 v11.5.4

Shaikh Quader, *AI Architect, IBM Db2*

LinkedIn: /in/shaikhquader

Medium: @shaikhquader

Twitter: @ShaikhQuader

66%

ML projects use Relational data

Source: The State of Data Science & Machine Learning 2017, Kaggle, October 2017
(based on 2017 Kaggle survey of 16,000 ML practitioners)

Machine Learning Workflow in an Enterprise

ML Dev Challenges:

- ❑ **Regulated data** (e.g., GDPR)
- ❑ **Data Volume** (data >> Dev env)
- ❑ **Noisy data** (25% data are noisy)

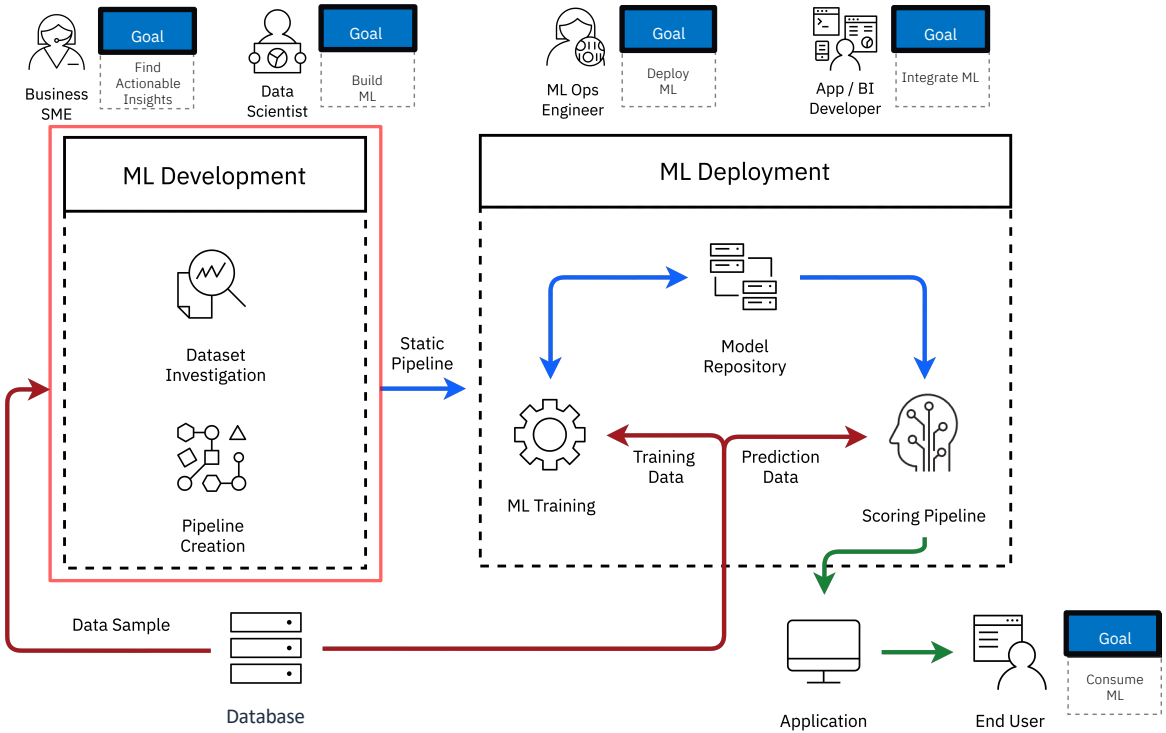

Machine Learning Workflow in an Enterprise

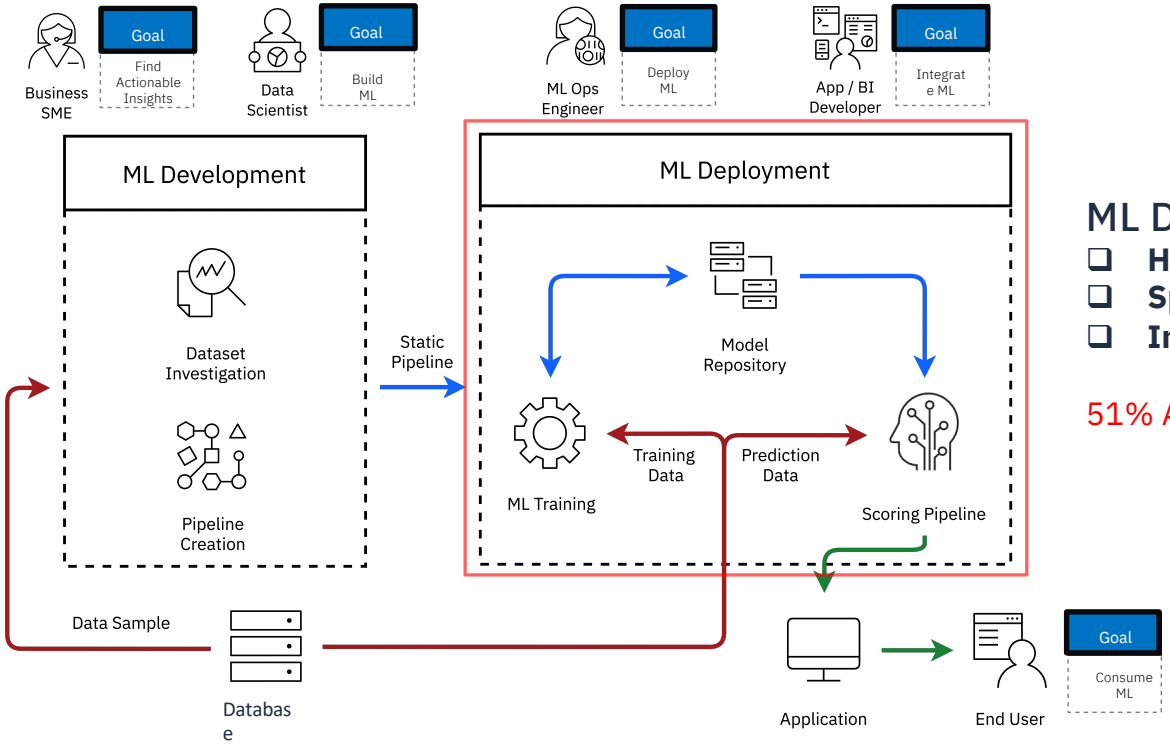

ML Deployment Challenges:

- ❑ **Hosting**
- ❑ **Speed**
- ❑ **Integration**

51% AI projects don't go beyond experiments

Accelerating and Optimizing AI Lifecycle with IBM Db2

01

Integrating Open-source Models with Db2

02

Developing and Deploying Db2-Native ML Models

Solution 1:
Bring your Open-Source Models to Db2

While the models may be centrally trained, the resulting inference pipelines will be deployed everywhere…to make inferences ("scoring") where the data is.

Agrawal et al.
Cloudy with high chance of DBMS: A 10-year prediction for Enterprise-Grade ML (CIDR '20)

Use Cases for in-database Scoring

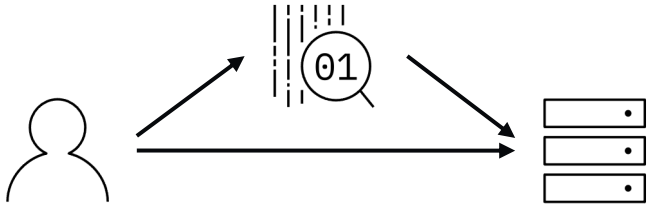

Latency-sensitive Decisions

Instantaneous predictions

Examples:
- Payment processing
- Fraud detection
- Loan/claim pre-approval

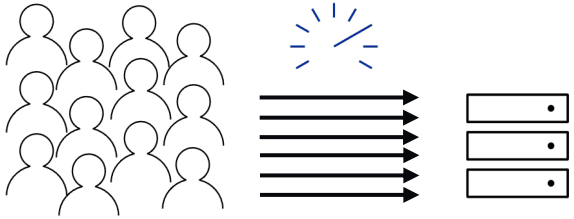

Large Batch Predictions

Real-time prediction using “fresh” and large operational data

Examples:
- Anomaly detection
- Escalation risk prediction
- Dynamic price optimization

Python UDF: Scoring Python Models via Db2

Watson
or
Open Source

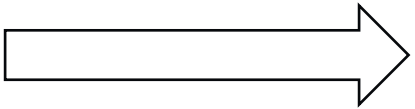

Export the ML pipeline by serializing *python joblib*

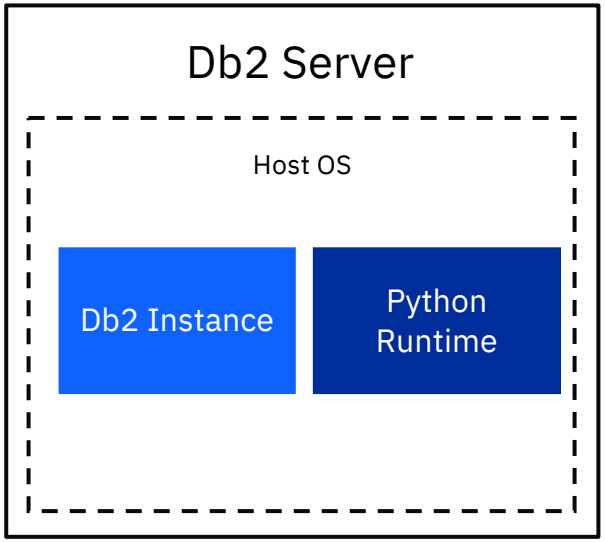

In-db Inferencing

Benefits:

- ML Infrastructure
- Low-latency
- High-throughput
- Simpler Integration

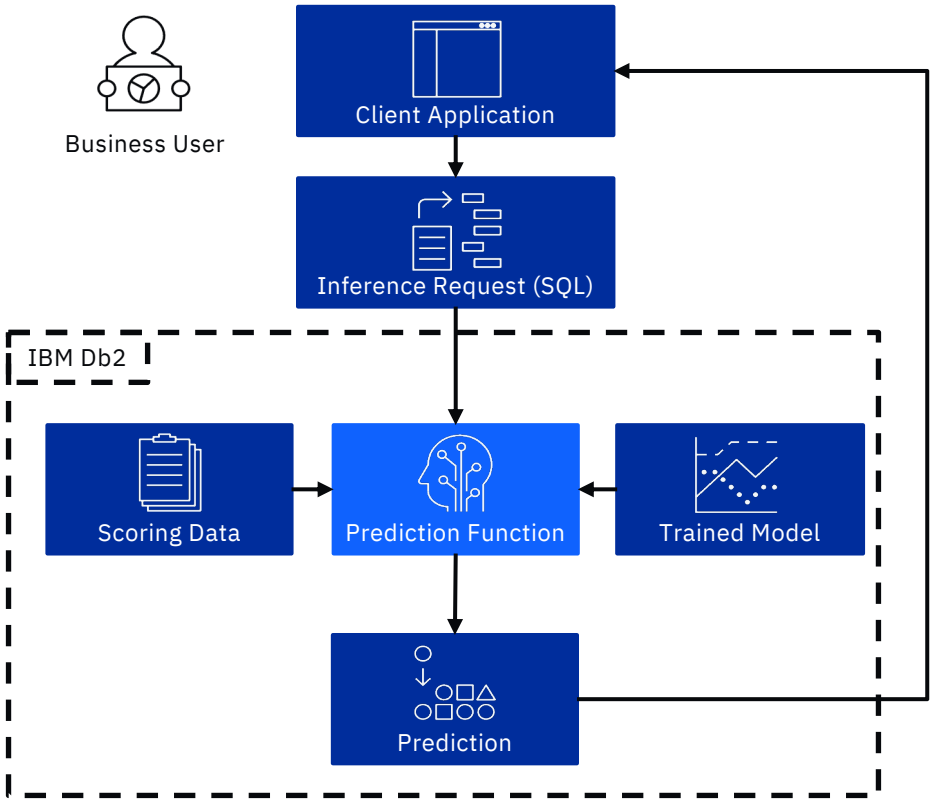

5x Speed up of Inferencing with IBM Db2

Model & Dataset

Logistic Regression (scikit-learn) Model

Batch size: 800k rows

Scoring on a Separate System

Data retrieved (over network) from Db2, scored, and written back to Db2: **1m47s**

Scoring with Db2

Data retrieved, scored, and written back inside Db2: **21s**

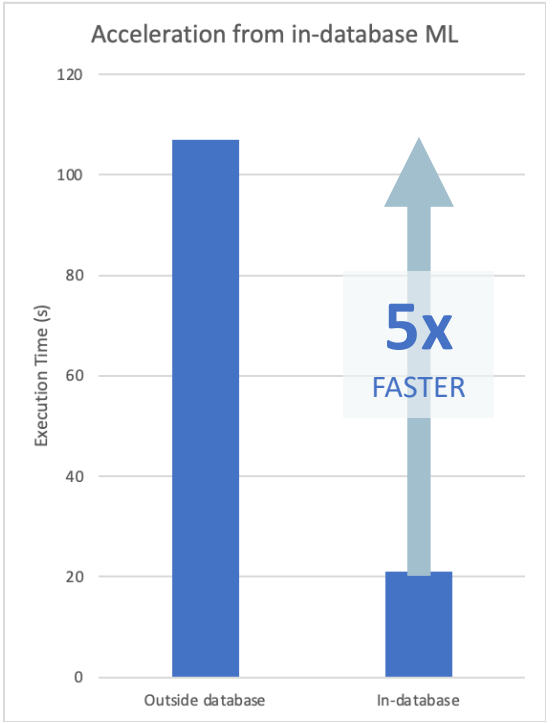

Solution 2:
Build and Deploy ML Models inside the Database

Db2 11.5.4 ML Capabilities

- Speed
- SQL Interface
- Distributed Training
- Security
- Scalability

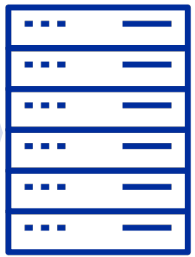

Platforms
x86 Linux
PPCLE
zLinux (Linear Reg. unsupported)

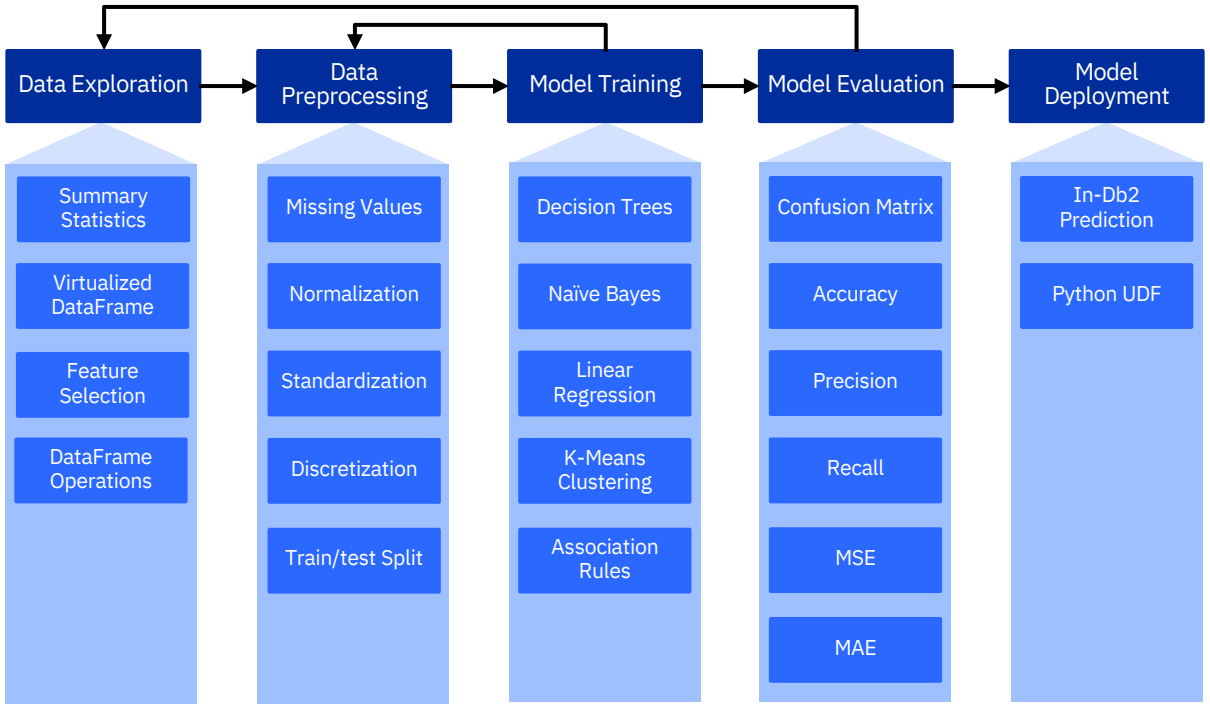

Distributed Model Training

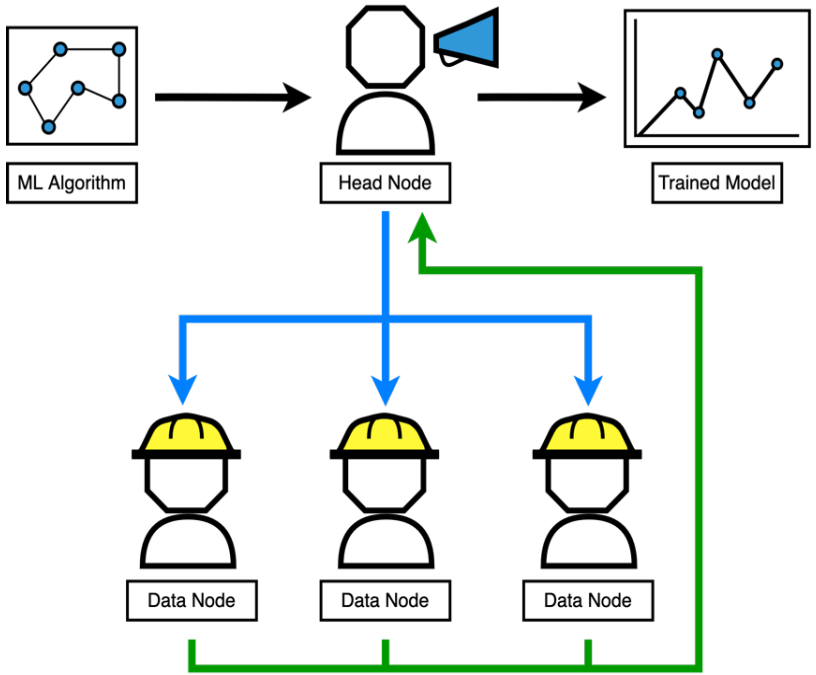

Real Use Case: Product Recommendations

IBM Sales team is finding product recommendations for customers using Db2's ML algorithms on the data stored in Db2 :

- Segmenting customers using k-means clustering algorithm
- Finding product recommendations using association rules mining algorithm

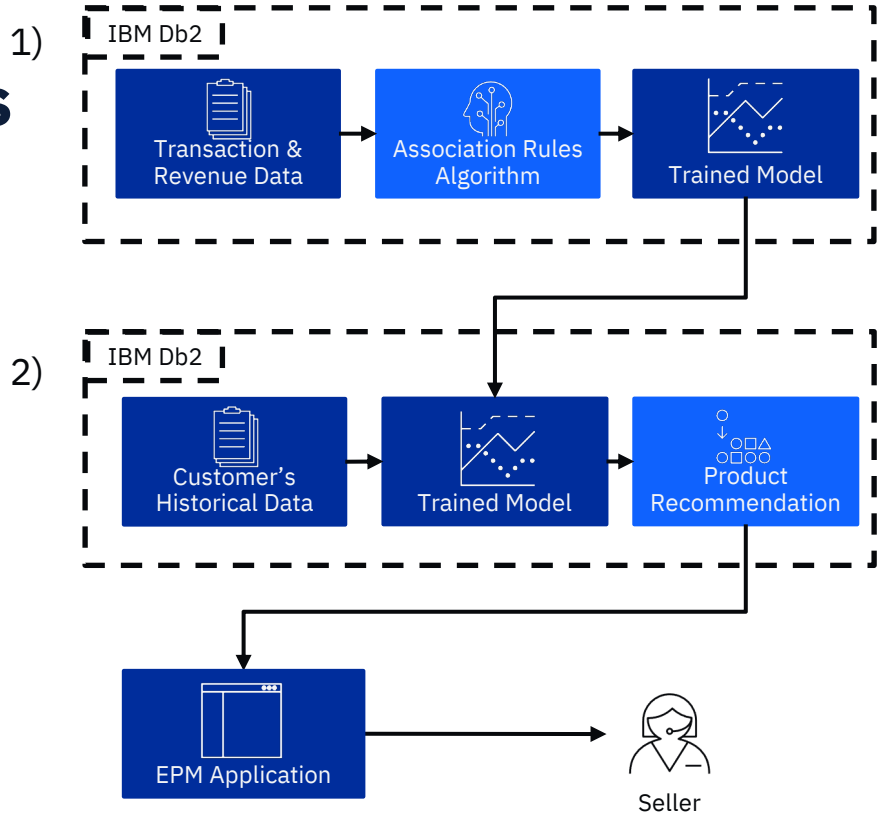

Demo:
Developing and Deploying Db2-Native ML Models

In-Db2 Model Building & Automated Model Deployment

Benefits:
- Speed
- IT cost savings
- Simpler Architecture
- Simplified Governance

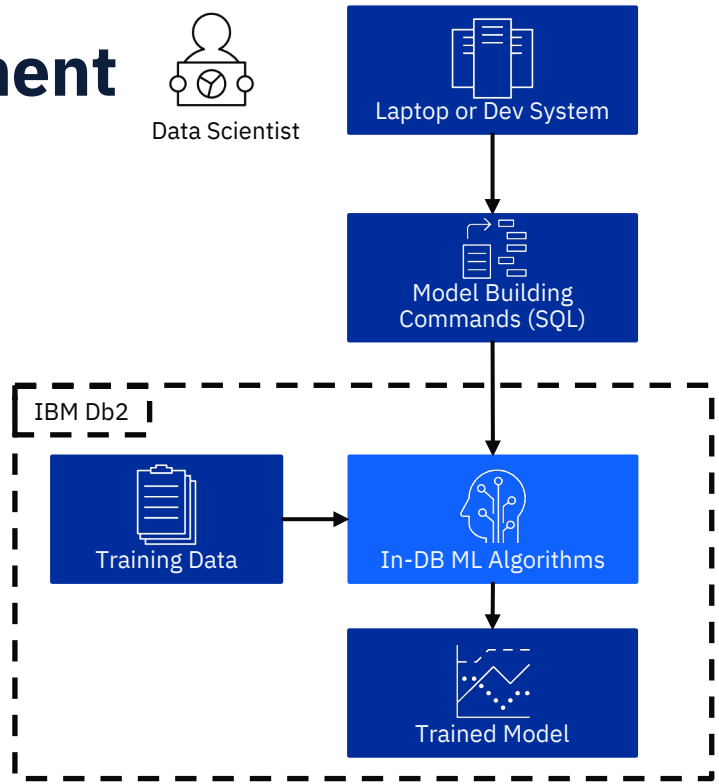

Challenges	Benefits of ML with Db2
Talent Gap	SQL interface for ML Deployment
Sensitive data	Virtualized data access and model building
Data transfer costs	In-Db2 ML – no data transfer
Infrastructure	Secure and scalable Db2 Infrastructure
Inference performance	In-Db2 scoring of Db2-native and open source models
Integration of ML	SQL-based Inference Queries

Demos and Tutorials

Demos:

- Build a Customer Segmentation Model with Db2 (K-Means Clustering)
- Build a Classification Model with Db2 (Decision Tree)
- Build a Regression Model with Db2 (Linear Regression)
- Integrate a Db2-native model with a Cognos Dashboard

Hands-On:

- Tutorials and Jupyter Notebooks

Documentation:

- Db2 11.5 Knowledge Center