



# Intro to AI and IBM i

## Mike Hollinger

Master Inventor, Principal Engineer PowerAI Vision  
IBM Cognitive Systems  
🐦@mike\_Hollinger

## Jesse Gorzinski


Business Architect  
IBM Cognitive Systems  
🐦@IBMJesseG

## Mark Irish

Software Developer  
IBM Cognitive Systems  
🐦@markdirish


## Alison Butteril

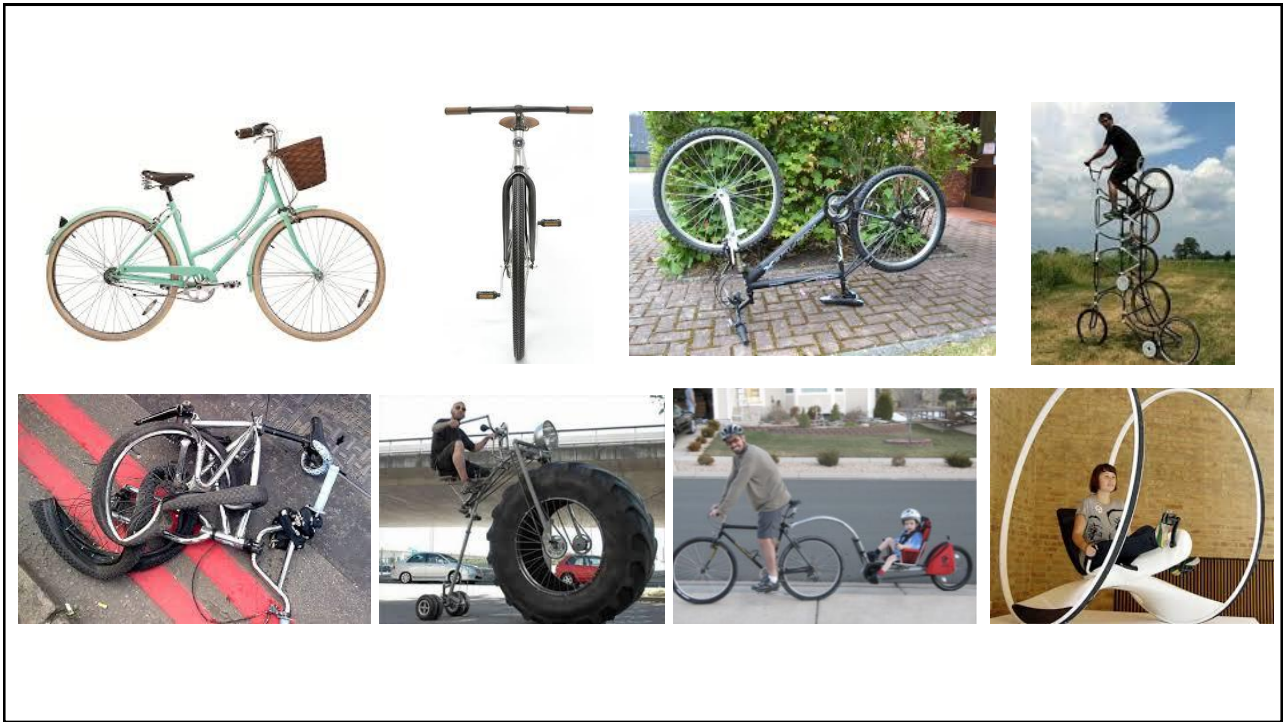
IBM i Offering Manager  
IBM Cognitive Systems  
🐦@IBMiSight



## 0. Warmup Exercise

# What's a bicycle look like?






2011



26% error rate

Machine Learning Based



5% error rate

Humans

2016

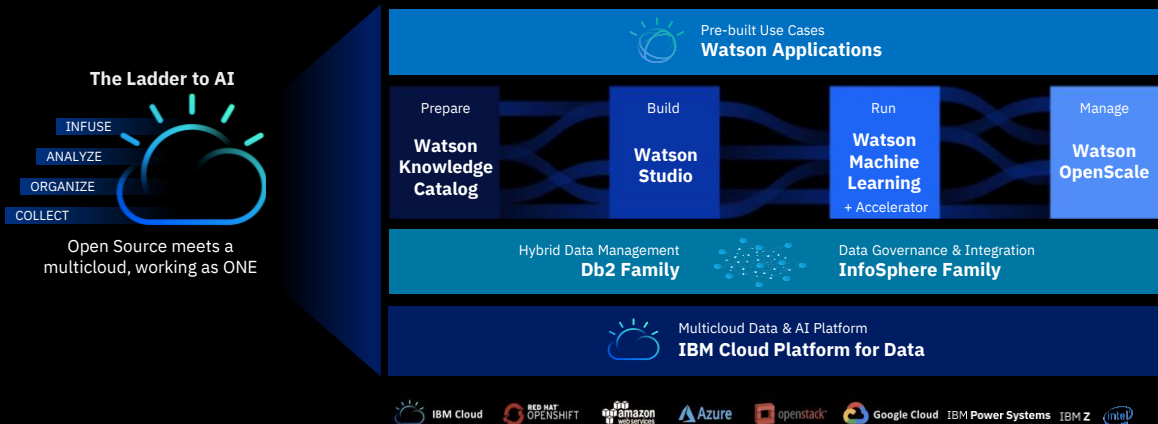


3% error rate

Deep Learning Based

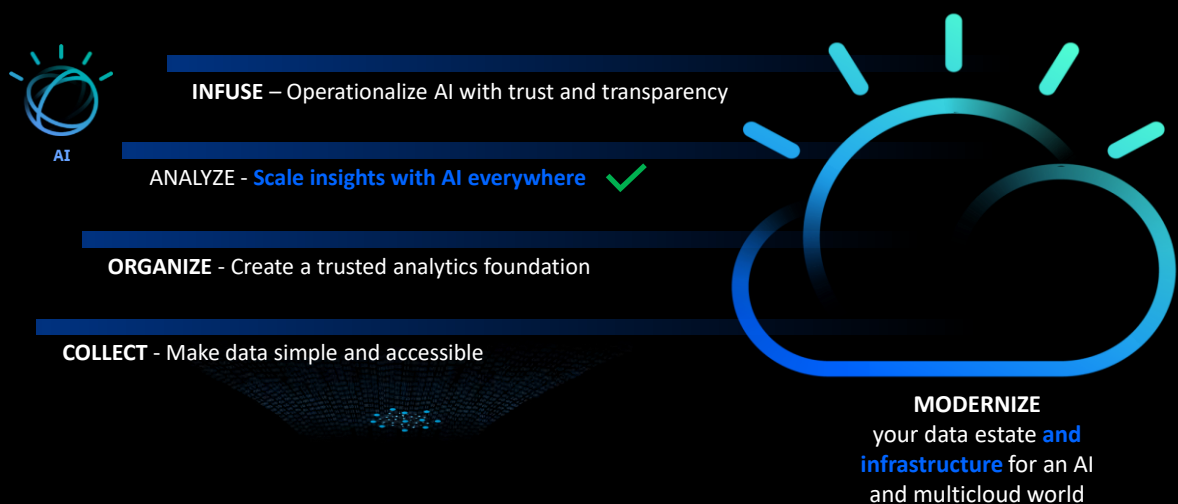
# IBM Data & AI Portfolio

Everything you need for Enterprise AI, on any cloud

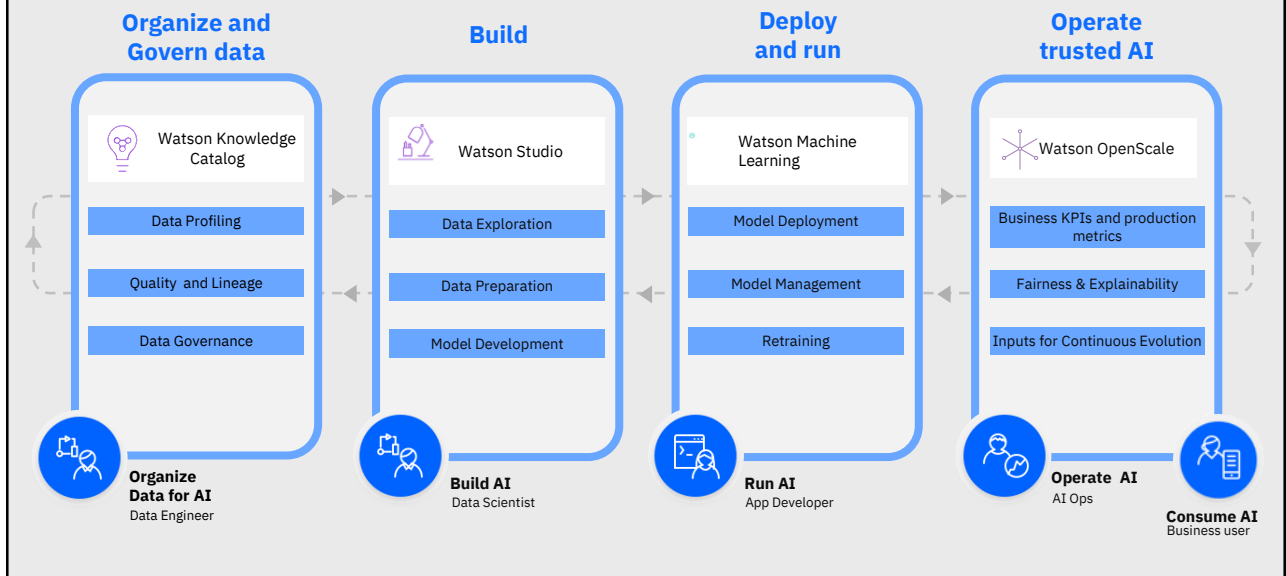


## The AI Ladder

A prescriptive approach to accelerating the journey to AI



## Watson Studio, WML and Watson OpenScale enable clients to operationalize AI across the enterprise



## OFFICIAL NOTICE



- **I am not a data scientist!**
- Data scientists spend **years learning how to analyze data**, I spent years learning how to develop software.
- I do not know all of the correct algorithms, models, or statistical tests to get the most out of your data.
- If you want a personalized understanding of your data, **you are going to have to hire a data scientist.**

# Machine Learning



- Machine learning is about training your computer on lots of data so it can learn how to do things you don't want to (or can't) do.
- Through machine learning, computers can learn things that we know, but just don't want to do.
  - Speech to text transcription
  - Language translation
- Through machine learning, computers can see relationships that are too complex for our brains to handle.
  - What factors indicate immediate failure of important components?
  - How can we identify credit accounts that are likely to default?

© 2019 IBM Corporation

# AI? ML? DL?



- There are three related terms used to describe machine learning:
  - **Artificial Intelligence (AI):** Computers mimicking human intelligence and make 'smart' decisions.
  - **Machine Learning (ML):** Algorithms and statistical techniques allowing computers to make smarter decisions with more data.
  - **Deep Learning (DL):** Computers training themselves on massive quantities of data using 'multilayered neural networks.'

© 2019 IBM Corporation


## AI? ML? DL?



- There are three related terms used to describe machine learning:
  - **Artificial Intelligence (AI):** Computers doing smart things.
  - **Machine Learning (ML):** Computers building a model with your help.
  - **Deep Learning (DL):** Computers building a model by themselves.

© 2019 IBM Corporation

## AI? ML? DL?



© 2019 IBM Corporation

## Training? Inferencing?



- **Training:** Taking your data, selecting an algorithm, and training a ‘model’ that can make educated predictions
- **Inferencing:** Drawing inferences between data points of a record to draw some conclusion.
- **A model is trained with data records that you have known results for. Your model inferences on new data with unknown outcomes it to try and deliver new insight.**

---

© 2019 IBM Corporation

## Fuzzy Logic, Stochastic Computing



- Machine learning is about getting ‘good enough’
- Like humans, machines are not 100% accurate (yet)
- Machine learning is about different business needs than your deterministic RPG, C, CL, and COBOL programs!

---

© 2019 IBM Corporation




## So how does it work?



### Automate



### Understand



One "RIGHT" Answer

Payroll Calculation

The "BEST" Answer

Influencing a decision

© 2019 IBM Corporation

## Approximate data



- Key requirements
  - approximation should only be in non-critical data
    - approximating critical data could lead to disastrous consequences
  - Identifying the section(s) of an application that could be approximated
    - Need application programmer & application domain expert involved



© 2019 IBM Corporation


## Modern Workloads



### DETERMINISTIC



100s of GBs of Data  
Noise Intolerant  
High Precision  
High Energy/Computation  
High Cost



### STOCHASTIC




10s of PBs of Data  
Noise tolerant  
Amenable to lower Precision  
Lower Energy/Computation  
Lower Cost

© 2019 IBM Corporation

## Machine Learning for Modern Workloads



© 2019 IBM Corporation



# Machine Learning with the IBM i

© 2019 IBM Corporation

## Why on the i?



- Machine learning requires **heaps of data** to be accurate.
- Your IBM i machines have been **accumulating data for years**, plenty of insights to be drawn
- **IBM is at the forefront** of a machine learning revolution
  - ~~Watson Services and APIs~~ **Watson Machine Learning Services** comes pre-trained, can be used with your data for data predictions.
  - **Watson Machine Learning Studio** allows you to easily create your own models in the cloud.
  - ~~PowerAI~~ **Watson Machine Learning Accelerator or Community Edition** is a suite of tools for your to train your own models on your own data with an AC922.

© 2019 IBM Corporation

## But...



- Training your ML/DL model requires a lot of simple matrix calculations.
- CPUs like the POWER processors in your IBM i machines are *OK*, but **pale in comparison to GPUs**.
- IBM i currently has no GPU support.




A GPU has more Arithmetic Logic Units (ALU) than a typical CPU.  
 • Increased ability to process simple operations in parallel

© 2019 IBM Corporation

## Watson Machine Learning Community Edition



- POWER chips using NVLink quickly send and receive data to GPUs
- No GPUs on IBM i means you must:
  - Buy Power Systems servers (or time on the server)
  - Send your data over to be quickly trained
  - Return the model to your IBM i (or send data to your server to do inferencing).



© 2019 IBM Corporation

## Popular ML Software



- Scikit Learn
- TensorFlow
- Caffe
- Pytorch
- Theano
- Spark



© 2019 IBM Corporation

## Popular ML Software on the IBM i



- Scikit Learn
- TensorFlow
- Caffe
- Pytorch
- Theano
- Spark



© 2019 IBM Corporation

## Some good news



- RPM enablement
- Python enablement
- BLAS enablement
- Db2 connection

Packages that enable ML are readily available

- Numpy, Pandas for data processing
- Scipy, Scikit Learn for ML and scientific analysis
- ipython, interactive python language support
- nltk, natural language toolkit for natural language ML process.
- matplotlib, jupyter notebook for visual/interactive ML/data analysis

© 2019 IBM Corporation

## Near future



The DL frameworks on our todo list

- Tensorflow
- Pytorch
- Caffe

Possible enhancements from low level of system

- Enhance BLAS or other math libraries to improve performance
- Expose more AI-helpful features from Power 9 or higher
- GPU support on IBM i

© 2019 IBM Corporation

## Popular ML Software on the IBM i (soon?)



- Scikit Learn
- TensorFlow
- Caffe
- Pytorch
- Theano
- Spark



© 2019 IBM Corporation



## IBM i Machine Learning Example

© 2019 IBM Corporation

## Machine Learning on IBM i example



Credit default prediction on a publicly available dataset:

- <https://github.com/elenalowery/DSX-Local-Credit-Card-Default/tree/master/data>
- 2 files: CUST\_HISTORY.csv, NEW\_CUSTOMERS.csv
  - CUST\_HISTORY.csv: Holds 26 data columns about 1000 customers, including **whether they defaulted**.
  - NEW\_CUSTOMERS.csv: Holds 25 data columns about 93 new customers (**don't know whether they defaulted**).

© 2019 IBM Corporation

## Machine Learning on IBM i example



### ▪ CUST\_HISTORY.csv

Q	R	S	T	U	V	W	X	Y	Z	AA
PROP_UNKN	ESTABLISHED_MONTH	OTHER_INSTALL_PLAN	RENT	OWN_RESIDENCE	NUMBER_CREDITS	RFM_SCORE	BRANCHES	TELEPHONE	SHIP_INTERNATIONAL	IS_DEFAULT
NO	38	YES	NO	YES	2	2	1	NO	YES	No
NO	34	NO	NO	YES	2	3	1	YES	NO	No
NO	29	NO	NO	YES	2	3	1	YES	NO	No
NO	31	YES	NO	YES	1	4	2	YES	NO	No
NO	28	NO	YES	NO	1	4	1	YES	NO	No
NO	35	NO	NO	YES	1	3	1	NO	YES	No
NO	33	YES	YES	NO	1	2	1	NO	NO	No
NO	42	NO	NO	YES	1	3	1	NO	NO	Yes
NO	43	YES	YES	NO	1	3	1	YES	NO	Yes
NO	44	NO	NO	YES	2	3	1	YES	NO	No
YES	42	NO	NO	NO	2	3	1	NO	NO	No
NO	40	NO	YES	NO	1	3	1	YES	NO	No
NO	36	NO	NO	YES	1	4	1	YES	NO	No
NO	20	NO	YES	NO	1	4	2	YES	NO	No
NO	24	NO	NO	YES	1	2	1	NO	YES	No
NO	27	NO	NO	YES	1	3	1	NO	NO	No
NO	46	NO	NO	YES	2	2	1	NO	NO	No
NO	33	NO	NO	YES	1	2	1	NO	NO	No
NO	34	NO	NO	YES	2	3	1	NO	NO	No
NO	25	YES	NO	YES	1	2	1	NO	NO	Yes
NO	25	NO	NO	YES	1	3	1	NO	NO	No
NO	28	NO	NO	YES	1	3	1	YES	NO	No
NO	31	NO	YES	NO	1	3	1	YES	NO	Yes
NO	32	YES	NO	YES	1	4	1	YES	YES	Yes
NO	32	NO	NO	YES	1	3	2	NO	NO	No
YES	68	YES	NO	YES	1	4	1	YES	NO	Yes

© 2019 IBM Corporation



# Machine Learning on IBM i example



## NEW\_CUSTOMERS.csv

Q	R	S	T	U	V	W	X	Y	Z	AA
PROP_UNKN	ESTABLISHED_MONTH	OTHER_INSTALL_PLAN	RENT	OWN_RESIDENCE	NUMBER_CREDITS	RFM_SCORE	BRANCHES	TELEPHONE	SHIP_INTERNATIONAL	
NO	49	NO	NO	YES	1	2	2	NO	NO	
NO	28	NO	NO	YES	2	4	1	NO	NO	
NO	48	NO	YES	NO	2	2	2	NO	YES	
NO	26	YES	NO	YES	1	2	2	NO	NO	
YES	52	NO	NO	YES	1	4	1	YES	NO	
NO	23	NO	NO	YES	1	4	1	NO	NO	
NO	46	NO	YES	NO	2	2	1	YES	NO	
NO	24	NO	YES	NO	1	2	1	NO	NO	
NO	52	NO	NO	YES	1	2	1	NO	NO	
NO	36	YES	YES	NO	2	4	2	YES	NO	
NO	30	NO	NO	YES	1	3	1	YES	YES	
NO	35	NO	NO	YES	1	3	1	YES	NO	
NO	28	NO	YES	NO	1	2	1	NO	NO	
NO	36	NO	NO	YES	1	4	1	YES	NO	
NO	36	NO	YES	NO	1	4	1	YES	NO	
NO	48	NO	NO	YES	2	3	2	NO	YES	
YES	70	YES	NO	NO	1	4	1	YES	NO	
NO	25	NO	YES	NO	1	3	1	YES	NO	
NO	68	NO	NO	NO	3	1	1	YES	NO	
NO	21	NO	YES	NO	1	3	1	NO	NO	
NO	24	NO	YES	NO	1	2	1	NO	NO	
YES	35	NO	NO	YES	1	3	1	NO	NO	
NO	42	NO	NO	YES	3	2	2	NO	NO	
NO	24	NO	NO	YES	1	3	1	YES	NO	
NO	55	NO	NO	YES	1	3	1	NO	NO	
NO	36	NO	NO	YES	2	3	1	NO	NO	
NO	25	NO	NO	YES	2	3	1	NO	NO	

© 2019 IBM Corporation

# Machine Learning on IBM i example



- You already have your data in Db2 on IBM i.
- So I import the data into the IBM i through Access Client Solutions (ACS).

The screenshot displays the IBM Access Client Solutions (ACS) interface. On the left, a sidebar shows the 'Data Transfer' section. The main window is titled 'Data Transfer' and shows a 'File View' of a local file system. A file named 'MARK.CUST.HISTORY' is selected. The 'Contents of MARK.CUST.HISTORY' window is open, showing a table with columns: MERCHANT, ACCT\_STATUS\_K\_USD, CONTRACT\_DURATION\_MONTH, HISTORY, CREDIT\_PROGRAM, and AMOUNT\_K\_USD. The table contains 100 rows of data, including various merchant names and account statuses. The bottom of the window shows '100 rows disclosed (more data available)' and buttons for 'Columns...', 'Save Results...', and 'Refresh'.

© 2019 IBM Corporation

## Machine Learning on IBM i example



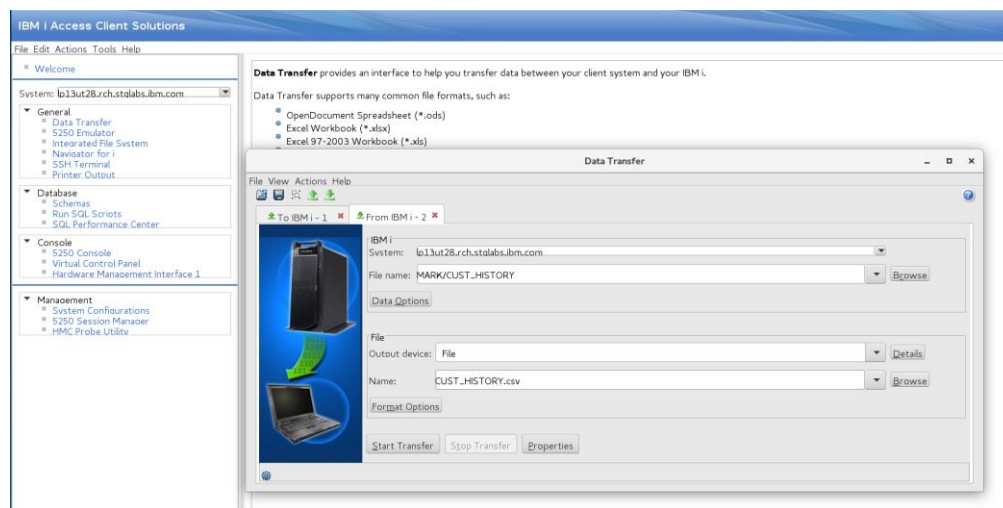
- First, we will use Watson Studio for Watson Machine Learning to build a model for our data.
- <https://www.ibm.com/cloud/garage/dte/tutorial/ibmr-watson-studio-ml-dl-made-easy>
- Instead of using 'customer\_churn.csv', use the data from your own table. In this case, from MARK/NEW\_CUSTOMER and MARK/CUST\_HISTORY.
- Need to get our data in .csv format... (this is where you might start)

© 2019 IBM Corporation

## Machine Learning on IBM i example



- ACS can be used to transfer data from IBM i to a .csv file.



© 2019 IBM Corporation

# Machine Learning on IBM i example



- With the .csv files of our data, we can import them into Watson Studio

The screenshot shows the IBM Watson Studio interface. The top navigation bar includes 'My Projects / Credit Default ML Example', 'Launch IDE', 'Add to project', and 'Upgrade'. The main content area is titled 'Data assets' and shows a table with two data assets: 'CUST\_HISTORY.csv' and 'NEW\_CUSTOMERS.csv'. Both are 'Data Asset' type, created by 'Mark Irish' on '9 Mar 2019, 10:11:06 am' and '9 Mar 2019, 10:11:05 am' respectively. A search bar at the top left asks 'What assets are you looking for?'. On the right, there is a 'Drop files here or browse for files to upload.' box.

NAME	TYPE	CREATED BY	LAST MODIFIED	ACTIONS
CUST_HISTORY.csv	Data Asset	Mark Irish	9 Mar 2019, 10:11:06 am	
NEW_CUSTOMERS.csv	Data Asset	Mark Irish	9 Mar 2019, 10:11:05 am	

© 2019 IBM Corporation

# Machine Learning on IBM i example



- We have our data, but no models...

The screenshot shows the IBM Watson Studio interface. The 'Data assets' section is expanded, showing the same two data assets as before. Below this, the 'Models' section is expanded, showing 'Watson Machine Learning models'. A button 'New Watson Machine Learning model' is visible. The table below is empty, indicating no models are present.

NAME	STATUS	TYPE	RUNTIME	LAST MODIFIED	ACTIONS
You don't have any Watson Machine Learning models yet.					

© 2019 IBM Corporation

# Machine Learning on IBM i example



- ...so we create a model, and have two different estimators...

### Select a technique

You cannot change label columns, feature columns, model type, or validation split after adding an estimator. You must first delete all estimators in order to make changes to these attributes.

Column value to predict (Label Col)

IS\_DEFAULT (String)

Feature columns

All (default)

**Binary Classification**

Classify new data into defined categories based on existing data. Choose if your label column contains two distinct categories.

**Multiclass Classification**

Classify new data into defined categories based on existing data. Choose if your label column contains a discrete number of categories.

**Regression**

Predict values from a continuous set of values. Choose if your label column contains a large number of values.

Validation Split

Train: 60      Test: 20      Holdout: 20

+ Add Estimators

### Configured estimators

Decision Tree Classifier

Not Yet Trained

Random Forest Classifier

Not Yet Trained

© 2019 IBM Corporation

# Machine Learning on IBM i example



- ...and see which model performs better.

### Select a technique

You cannot change label columns, feature columns, model type, or validation split after adding an estimator. You must first delete all estimators in order to make changes to these attributes.

Column value to predict (Label Col)

IS\_DEFAULT (String)

Feature columns

All (default)

**Binary Classification**

Classify new data into defined categories based on existing data. Choose if your label column contains two distinct categories.

**Multiclass Classification**

Classify new data into defined categories based on existing data. Choose if your label column contains a discrete number of categories.

**Regression**

Predict values from a continuous set of values. Choose if your label column contains a large number of values.

Validation Split

Train: 60      Test: 20      Holdout: 20

+ Add Estimators

### Configured estimators

Decision Tree Classifier

Not Yet Trained

Random Forest Classifier

Not Yet Trained

© 2019 IBM Corporation

## Machine Learning on IBM i example



- ...and see which model performs better.

### Select model

	ESTIMATOR TYPE	STATUS	PERFORMANCE	AREA UNDER ROC CURVE	AREA UNDER PR CURVE
<input checked="" type="radio"/>	RandomForestClassifier	Trained & Evaluated	Fair	0.76497	0.46623
<input type="radio"/>	DecisionTreeClassifier	Trained & Evaluated	Fair	0.70093	0.36183

© 2019 IBM Corporation

## Machine Learning on IBM i example



- We save our model, then deploy it as a web service.

### Credit Default Analysis

Overview Evaluation Deployments Lineage

NAME	STATUS	DEPLOYMENT TYPE
credit	DEPLOY_SUCCESS	Web Service

© 2019 IBM Corporation

# Machine Learning on IBM i example



- Watson Studio will tell you how to send data to your service in a number of programming languages and utilities.

credit

Overview Implementation Test

**Implementation**

Scoring End-point [https://us-south.ml.cloud.ibm.com/v3/ml\\_instances/775792de-777c-4cc3-9daa-43c706662ba/deployments/cead1d29-9a63-41cd-86ff-0a57b066e81d/online](https://us-south.ml.cloud.ibm.com/v3/ml_instances/775792de-777c-4cc3-9daa-43c706662ba/deployments/cead1d29-9a63-41cd-86ff-0a57b066e81d/online)

Authorization: Bearer <token> See code snippets below for information on how to retrieve the IBM Authorization Token to be passed with scoring requests.

Content-type: application/json Required if the request body is sent in JSON format.

**Code Snippets**

cURL Java JavaScript Python Scala

```
const XMLHttpRequest = require("xmlhttprequest").XMLHttpRequest;
const btoa = require("btoa");
const wml_credentials = new Map();

// NOTE: you must manually construct wml_credentials hash map below using information retrieved
// from your IBM Cloud Watson Machine Learning Service Instance

wml_credentials.set('url', wml_service_credentials.url);
wml_credentials.set('username', wml_service_credentials.username);
wml_credentials.set('password', wml_service_credentials.password);

function apiGet(url, username, password, loadCallback, errorCallback){
  const oReq = new XMLHttpRequest();
  const tokenHeader = "Basic " + btoa(username + ":" + password);
  const tokenUrl = url + "/v3/identity/token";
```

© 2019 IBM Corporation

# Machine Learning on IBM i example



- Results from Node.js:

```
[Array],
0,
'No',
[Array] ] ] }
[ 26,
[ 0, 2, 3, 4, 5, 7, 8, 14, 15, 17, 21, 22, 23 ],
[ 0.003586656685545,
0.9776464063731196,
0.9145825466646116,
3.4333491482067946,
0.7152380740805903,
2.331738584515539,
0.9159623110504432,
1.96907856389488,
2.209072203437452,
4.196511861662977,
1.7129974283616016,
3.1235357358694746,
5.7705132812785145 ] ]
[ 16.06982302958149, 3.9301769704185125 ]
[ 0.8034911514790745, 0.19650884852092562 ]
ml-example]$
```


© 2019 IBM Corporation

## Machine Learning on IBM i example



- Watson Studio will also let you test your model out.

credit



© 2019 IBM Corporation

## Machine Learning on IBM i example



- Credit default prediction on a publicly available dataset:
  - <https://github.com/elenalowery/DSX-Local-Credit-Card-Default/tree/master/data>
  - 2 files: CUST\_HISTORY.csv, NEW\_CUSTOMERS.csv
    - CUST\_HISTORY.csv: Holds 26 data columns about 1000 customers, including **whether they defaulted**.
    - NEW\_CUSTOMERS.csv: Holds 25 data columns about 93 new customers (**don't know whether they defaulted**).

© 2019 IBM Corporation

## Machine Learning on IBM i example



- Same thing can be done by hand!
- Developed with **scikit-learn**
- Training on **IBM I or Linux (Snap ML)**, with **inferencing directly on the IBM i**
- Uses **REST API in Node.js** to transfer data between Linux and IBM I
- Front-end web UI for demonstration (written with **Flask**)



- Note: Still finding a landing page for documentation and example source code. If you are interested, e-mail [mirish@ibm.com](mailto:mirish@ibm.com)

© 2019 IBM Corporation

## Machine Learning on IBM i example



Have 1000 known credit accounts. Train a model on 900 of them, then test model on the last 100 to determine model accuracy

0 = No Default, 1 = Default

Chance of No Default

Chance of Default

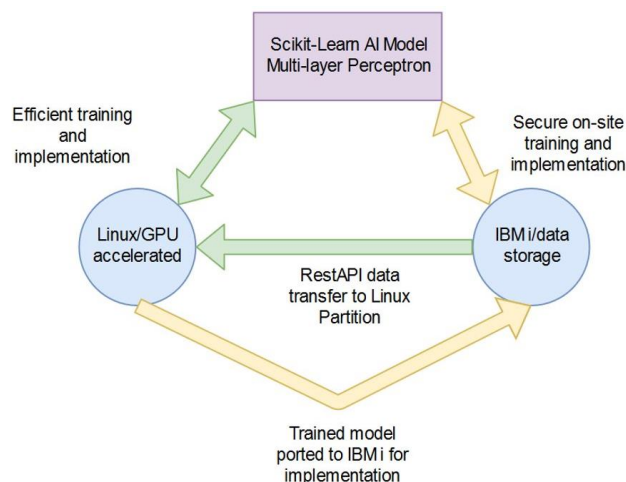
```
Prediction: 1 Actual: 0 Probability: ( 0.39931058938110997 , 0.60068941061889 )
Prediction: 1 Actual: 1 Probability: ( 0.04396396724667273 , 0.9560360327533273 )
Prediction: 0 Actual: 1 Probability: ( 0.8267071417589261 , 0.17329285824107388 )
Prediction: 0 Actual: 0 Probability: ( 0.9813813065900909 , 0.01861869340990902 )
Prediction: 0 Actual: 0 Probability: ( 0.7665148240541815 , 0.23348517594581847 )
Prediction: 0 Actual: 0 Probability: ( 0.9913802276766168 , 0.008619772323383209 )
Prediction: 0 Actual: 0 Probability: ( 0.747153939982939 , 0.252846060017061 )
Prediction: 0 Actual: 0 Probability: ( 0.996098291035753 , 0.003901708964246974 )
Prediction: 0 Actual: 1 Probability: ( 0.5766529003502558 , 0.4233470996497442 )
Prediction: 1 Actual: 0 Probability: ( 0.4769135583636863 , 0.5230864416363137 )
Prediction: 1 Actual: 0 Probability: ( 0.2410664026582301 , 0.7589335973417699 )
Prediction: 0 Actual: 0 Probability: ( 0.6893280492796938 , 0.3106719507203061 )
Prediction: 0 Actual: 0 Probability: ( 0.9750926365935764 , 0.024907363406423657 )
Prediction: 0 Actual: 0 Probability: ( 0.7565012828390749 , 0.24349871716092514 )
Prediction: 1 Actual: 0 Probability: ( 0.15274769798760635 , 0.8472523020123937 )
Prediction: 0 Actual: 1 Probability: ( 0.9805121676881574 , 0.019487832311842587 )
Prediction: 1 Actual: 1 Probability: ( 0.1434059901966399 , 0.8565940098033601 )
Prediction: 1 Actual: 1 Probability: ( 0.14677280883158828 , 0.8532271911684117 )
Prediction: 0 Actual: 0 Probability: ( 0.9774938908491175 , 0.022506109150882463 )
Accuracy of Neural Network classifier on test set: 0.69
```

What number of the 100 known accounts did the model correctly predict?

© 2019 IBM Corporation



# Machine Learning on IBM i example



© 2019 IBM Corporation

# Machine Learning on IBM i example



```

def main():
    # Connect to Db2 and get the data
    conn = db2.connect()
    cur = conn.cursor()
    cur.execute("SELECT * FROM MARK_NEW_CUST")

    rows = []

    # TODO: This is to convert to JSON, would help with transferring data because we could reuse
    class CursorByLine():
        def __init__(self, cursor):
            self._cursor = cursor

        def __iter__(self):
            return self

        def __next__(self):
            row = self._cursor.__next__()
            return {description[i]: row[col] for col, description in enumerate(self._cursor.description)}

    for row in CursorByLine(cur):
        rows.append(row)

    json_data = json.dumps(rows)

    model_path = './default_prediction_model.pkl'
    prepros_path = './feature_transform_model.pkl'
    model, le, scaler = load_model(model_path, prepros_path)
    processed_data, _ = process_file(json_data, le, scaler)
    predict_defaults(model, processed_data, rows)
  
```

Connect to Db2 to get data into Python, import the model, then inference on data with unknown outcomes.

```

-bash-4.4$ python predict.py
Gold Acoustics      Not likely to default ( 0.013 )
Wood Corp           Likely to default    ( 0.8644 )
Bridge Acoustics    Not likely to default ( 0.013 )
Rabbitechnologies   Might default        ( 0.2911 )
RhinoTainment       Not likely to default ( 0.1454 )
Voyagetrionics      Not likely to default ( 0.04 )
Webrows             Not likely to default ( 0.0476 )
Silversun           Not likely to default ( 0.0081 )
Redware             Not likely to default ( 0.1178 )
Leopardworth        Not likely to default ( 0.0225 )
Flux Networks       Likely to default    ( 0.6415 )
Storm Records       Not likely to default ( 0.0054 )
Petal Sports        Not likely to default ( 0.247 )
Midnightelligence   Not likely to default ( 0.0315 )
Pumpkinavigation     Not likely to default ( 0.1053 )
Canics              Not likely to default ( 0.0021 )
Shadowworks         Not likely to default ( 0.0473 )
Squidmart           Not likely to default ( 0.0699 )
Wondercoms          Not likely to default ( 0.0287 )
Arcanesun           Likely to default    ( 0.6612 )
Turtle Enterprises   Might default        ( 0.5432 )
Goblin Acoustics    Might default        ( 0.2554 )
Frostfire Media     Not likely to default ( 0.084 )
Phenomenologies     Not likely to default ( 0.1106 )
Pixsystems          Not likely to default ( 0.1683 )
Aces                Not likely to default ( 0.0297 )
Sphinxsecurity      Not likely to default ( 0.1738 )
Spikeshade          Likely to default    ( 0.86 )
Soulbeat            Not likely to default ( 0.1711 )
Bansheesys          Not likely to default ( 0.1485 )
Oyster Co.          Likely to default    ( 0.681 )
Soul                Not likely to default ( 0.1261 )
Wave Acoustics      Not likely to default ( 0.0771 )
Spectertainment     Not likely to default ( 0.1118 )
  
```

For each entry in the database, run against the model and print (and color) the output

© 2019 IBM Corporation

# Machine Learning on IBM i example



## Model Training

Take data with a known outcome and train a predictive model that determines how outcomes covary with other data. For this example, the data is credit account information, such as income, size of loan, and credit history, and the outcome is whether the account defaulted.

### Train from Db2 data

Use data from local Db2 table to train a model in scikit-learn locally.

Train

### Use pre-built model

Use a model that was previously built, either on the IBM i or from other hardware such as GPU-accelerated training on Linux or AIX.

Load

## Inferencing

New data can be run against the trained model, offering new and interesting insights that can be used in real business cases. In this example, enter data for a person seeking credit, and the model will attempt to determine whether they will default on their loan.

### Inference on new data

Account balance in thousands of USD

Choose...

Contract Duration in Months

CONTRACT\_DURATION\_MONTH

Credit History

Choose...

Credit Program

Choose...

Amount in thousands of USD

AMOUNT\_K\_USD

Account Type

Choose...

Account Age

Choose...

Present Residents

## Chance to Default

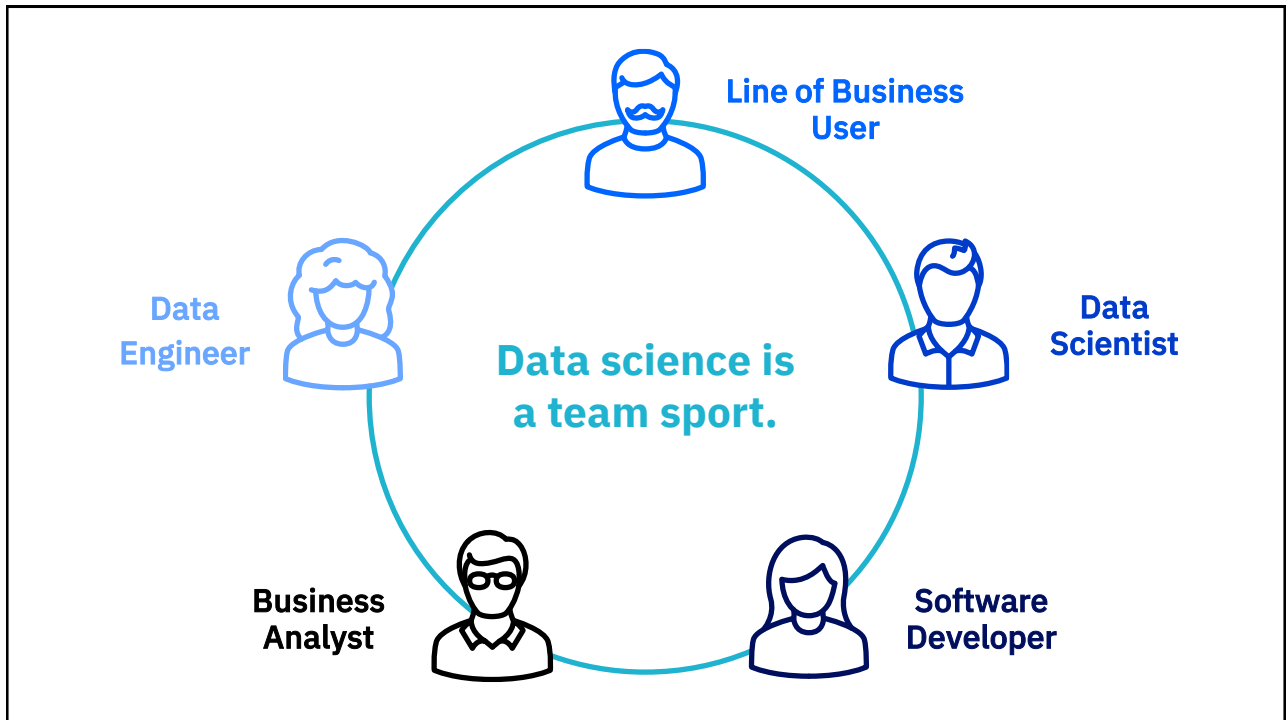
89.156%

Accept

Reject

© 2019 IBM Corporation

# PowerAI Vision: Advertising



IBM PowerAI Vision   Data Sets   Models   Custom Models   Deployed Models   admin   Log out

Welcome to IBM PowerAI Vision

**Create Dataset**

Start by adding images and video files to a data set.

**Prepare Data**

Label objects or assign categories to images or videos, then use auto labeling to complete the entire data set.

**Train Model**

Select a few custom options to create your model.

**Deploy Model**

Deploy the trained model and receive an API link for an inference device.

[Get started](#)

IBM PowerAI Vision Data Sets Models Custom Models Deployed Models admin Log out

### Data set / F1 Ads Dataset

Total files: 210 Matching files: 5 Selected files: 0

[Train model](#) [Augment data](#) [Auto label](#) [Export data set](#) View: grid list search

[Assign category](#) [Label objects](#) ☐ Select ☐ Delete [Refresh](#) Sort by: Select

Drop files here

Import files

Uncategorized

f1\_ferrari f1\_mercedes  
f1\_redbull ad\_rolex  
f1\_renault (13 more)

77 frames

Uncategorized

f1\_mercedes ad\_rolex  
ad\_pirelli f1\_ferrari  
f1\_force\_india (10 more)

38 frames

Uncategorized

ad\_johnnywalker f1\_ferrari  
f1\_mercedes ad\_rolex  
f1\_redbull (10 more)

46 frames

Uncategorized

No objects added

0 frames

Uncategorized

f1\_mercedes ad\_rolex  
f1\_redbull f1\_ferrari  
ad\_pirelli (7 more)

44 frames

Items per page: 20 | 1-5 of 5 items 1 of 1 pages < 1 >

IBM PowerAI Vision Data Sets Models Custom Models Deployed Models admin Log out

### Edit / USGP\_DEMO

[Done editing](#)

CATEGORY [Add New](#)

OBJECTS [Add New](#)

ad\_pirelli (7) ☐ f1\_mercedes (1) ☐ f1\_redbull (1) ☐ ad\_johnnywalker (0) f1\_unknown (0)  
ad\_komen (0) f1\_sauber (0) f1\_renault (0)

[+10 more](#)

VIDEO

Lap 2 / 56

1	VET	-1:31.2
2	HAM	-1:31.2
3	BOI	-1:32.8
4	RIC	-1:33.8
5	OCO	-1:35.2
6	RAI	-1:35.5
7	ALO	-1:35.8
8	STR	-1:37.5
9	MAS	-1:37.7
10	PER	-1:38.2
11	KVY	-1:38.4
12	GRO	-1:38.5
13	VIR	-1:38.8
14	ERI	-1:39.3
15	STR	-1:40.2
16	HIS	-1:40.9
17	VAN	-1:41.5
18	MAG	-1:41.5
19	WET	-1:41.5
20	HAR	-1:42.8

1:07 / 6:21

[Capture frame](#) [Auto capture frames](#) [Auto label](#)

FRAME

Lap 2 / 56

ad\_pirelli ad\_pirelli ad\_pirelli ad\_pirelli ad\_pirelli ad\_pirelli

f1\_mercedes f1\_redbull

35%

IBM PowerAI Vision    Data Sets    **Models**    Custom Models    Deployed Models    admin    Log out

Data set  
**USGP\_DEMO**

Objects

## Trained model / USGP\_DEMO\_model

Trained models are created from prepared data sets.  
The model can be validated, exported, and deployed for production for Graphics Processing Units (GPU).

**Object detection**      **Created**  
Optimized for accuracy (faster R-CNN)      12/5/2018, 2:24:50 PM

[Deploy model](#)    [Export model](#)      Advanced metrics: Off ☐ On

OBJECT	AVERAGE PRECISION	RECALL	IOU
<b>f1_ferrari</b> 111 objects / 95 images	0.892	0.6	0.892
<b>ad_dhl</b> 69 objects / 18 images	0.394	0.125	0.394
<b>f1_unknown</b> 84 objects / 31 images	0.005	0	0
<b>f1_redbull</b> 102 objects / 96 images	0.792	0.462	0.792
<b>ad_emirates</b> 105 objects / 19 images	0.027	0.065	0.027

IBM PowerAI Vision    Data Sets    Models    Custom Models    **Deployed Models**    admin    Log out

Data set  
**USGP\_Autolabel**

Objects

API endpoint: `api/dlapis/efd9b10c-a1f9-45e3-9df7-18c3f5445a87`    [Copy](#)    [GET](#)    [POST](#)

Test Model

Drop files here

[Import files](#)

To test a video using the API endpoint, click open.

[Open](#)

External URL

URL

[Upload](#)


## Results

OBJECTS	RESULT	AVERAGE
> <b>f1_mercedes</b> 2 objects	-	0.991
> <b>f1_force_india</b> 2 objects	-	0.934
<b>f1_ferrari</b> 1 objects	0.980	-
<b>f1_redbull</b> 1 objects	0.985	-



# PowerAI Vision: Retail Analytics







## Conclusions

© 2019 IBM Corporation



## Machine Learning is about fuzzy logic

- Algorithms can tell you interesting things about your data without you writing custom code.
- Instead of writing code, you feed the algorithms data to build models.
- These models can make (very well-)educated guesses to draw new insights on your data.
- Machine learning can automate things that previously could not be automated
- Machine learning can also find relationships too complex for us to see

© 2019 IBM Corporation



## Machine Learning is the present



- In the Information Age, data is growing faster than we realize
- A lot of the data is 'uncertain', in that we have data points but not the complete picture
- Machine learning helps draw conclusions about this uncertain data
- The most successful companies of the future will be those that utilize machine learning **NOW** to get ahead of the competition.

---

© 2019 IBM Corporation

## Machine Learning is possible on IBM i



- There are **many avenues for doing machine learning** with the data on your IBM i.
  - Watson Machine Learning (**Easiest**)
  - Watson Machine Learning Community Edition (**Harder, Fast**)
  - Open-source Python packages on IBM i (**Harder, Slow**)
- You have options!

---

© 2019 IBM Corporation

## Machine Learning already bringing value



See <http://ibm.biz/ibmistories>

- Computer Merchants (B2B Computing, Australia)
  - Using custom app which leverages AI to “identify issues that most likely require an engineer’s expertise”
- Vision Banco (Finance/Banking, Paraguay)
  - H2O Driverless AI on IBM Power Systems AC922 running Red Hat Enterprise Linux. The bank is now migrating all of their models to H2O Driverless AI. They are able to get insights into their data much faster than on traditional x86 infrastructure.
  - Saved time and increased revenue by building and deploying models that have doubled the number of credit products accepted per customer. Additionally, the team was able to pinpoint credit and default risks with greater accuracy than previously.

© 2019 IBM Corporation

## Machine Learning already bringing value



See <http://ibm.biz/ibmistories>

- Jori (Furniture, Belgium)
  - Leveraging scikit-learn to predict where to build retail locations. Also doing language translation.
- Robertet (Fragrance & Flavors, France)
  - Running chatbots based on Java and Watson
- Caixa Geral de Depositos (Finance/Banking, France)
  - A hybrid cloud credit-scoring application that uses advanced machine learning with IBM Watson to rapidly and accurately determine each customer’s credit-worthiness. Consumed via simple API call

© 2019 IBM Corporation

## Resources to help you get started



- IBM i OSS examples github
  - Tutorials and sample code for machine learning (and python, Node.js, ...)
  - <https://github.com/ibm/ibmi-oss-examples>
- Watson Machine Learning Demos
  - Available online, and so easy to use!
  - <https://www.ibm.com/cloud/garage/dte/tutorial/ibmr-watson-studio-mldl-made-easy>
  - There are free subscriptions that allow you to try before you buy!

© 2019 IBM Corporation

## Leaders everywhere are monetizing data & developing strategies to embed AI in business



### Retail

Market basket analysis, [Next best offer](#), [Customer churn](#), Propensity to buy




### Marketing

Email optimization, Lifetime client value, [Discount targeting](#)



### Healthcare

Medicare fraud, AI-assisted diagnosis, Drug demand forecast, [Medical Imaging](#)



### Manufacturing

Predictive maintenance, Process optimization, Demand forecast, [Product inspection & quality](#)



### Energy and Utilities

Power usage prediction, Smart grid management, [Asset inspection](#)



### Banking

Customer segmentation, Credit risk, credit card fraud detection, [Market prediction](#), ATM fraud



### Security

Malicious activity detection, Logs analysis, [Fraud analysis](#), [Video Analysis](#)



### Travel & transportation

Dynamic pricing, Call center assistants, Tourism forecasting, [Self-driving vehicles](#)

© 2019 IBM Corporation

Imagine the  
**possibilities** for deep  
learning



Imagine the  
**possibilities** for deep  
learning

...and for your  
business.



Let's put smart to work.™

