



AI TECH MCR-2

AI User Group

10th Sep 2018

#AIUserGroup @aitechUG

Guest speaker: Fahd Rafi, Microsoft

Host: Sherin Mathew

Partnered with



Sponsored by



Affiliated By




Agenda



- Registration and drinks 6:00 pm
- Opening - 6:15 pm
- Introduction - Sherin Mathew - 6:20 pm
 - ML Introduction and demo later (10 min max)
- ML with Spark on Databricks- Fahd Rafi- 6:30pm
- AI Demo - Fahd - 7:00 pm (20 max)
- AI Forum - Open discussion 7:20pm
 - Questions and debate - Three questions - 5 mins each
- Break - Food and drinks 7:35pm
- Networking Open forum - Demos and Networking - 8:00pm

About me

- Principal Consultant @Simpsons , Microsoft Gold Partner
- Azure Architecture - Enterprise BI Solution
- Passionate Data architect and “Azure Addict”
- MCP, MSc and B.Eng. MSP, Agile, Prince2
- Sports: Cricket
- Hobbies: Dancing, Recycled-Art
- Learning: To fly Helicopter – G2 Cabri
- Other interests: Classic Sports car
- Twitter: @smdisrupt



AI User Group



- Our Purpose
- Our Goals
- Announcements
 - Next Event
 - Entry on the website
- Roadmap –Partners, Speakers, Bloggers
- Feedback
- #AIUserGroup @aitechug



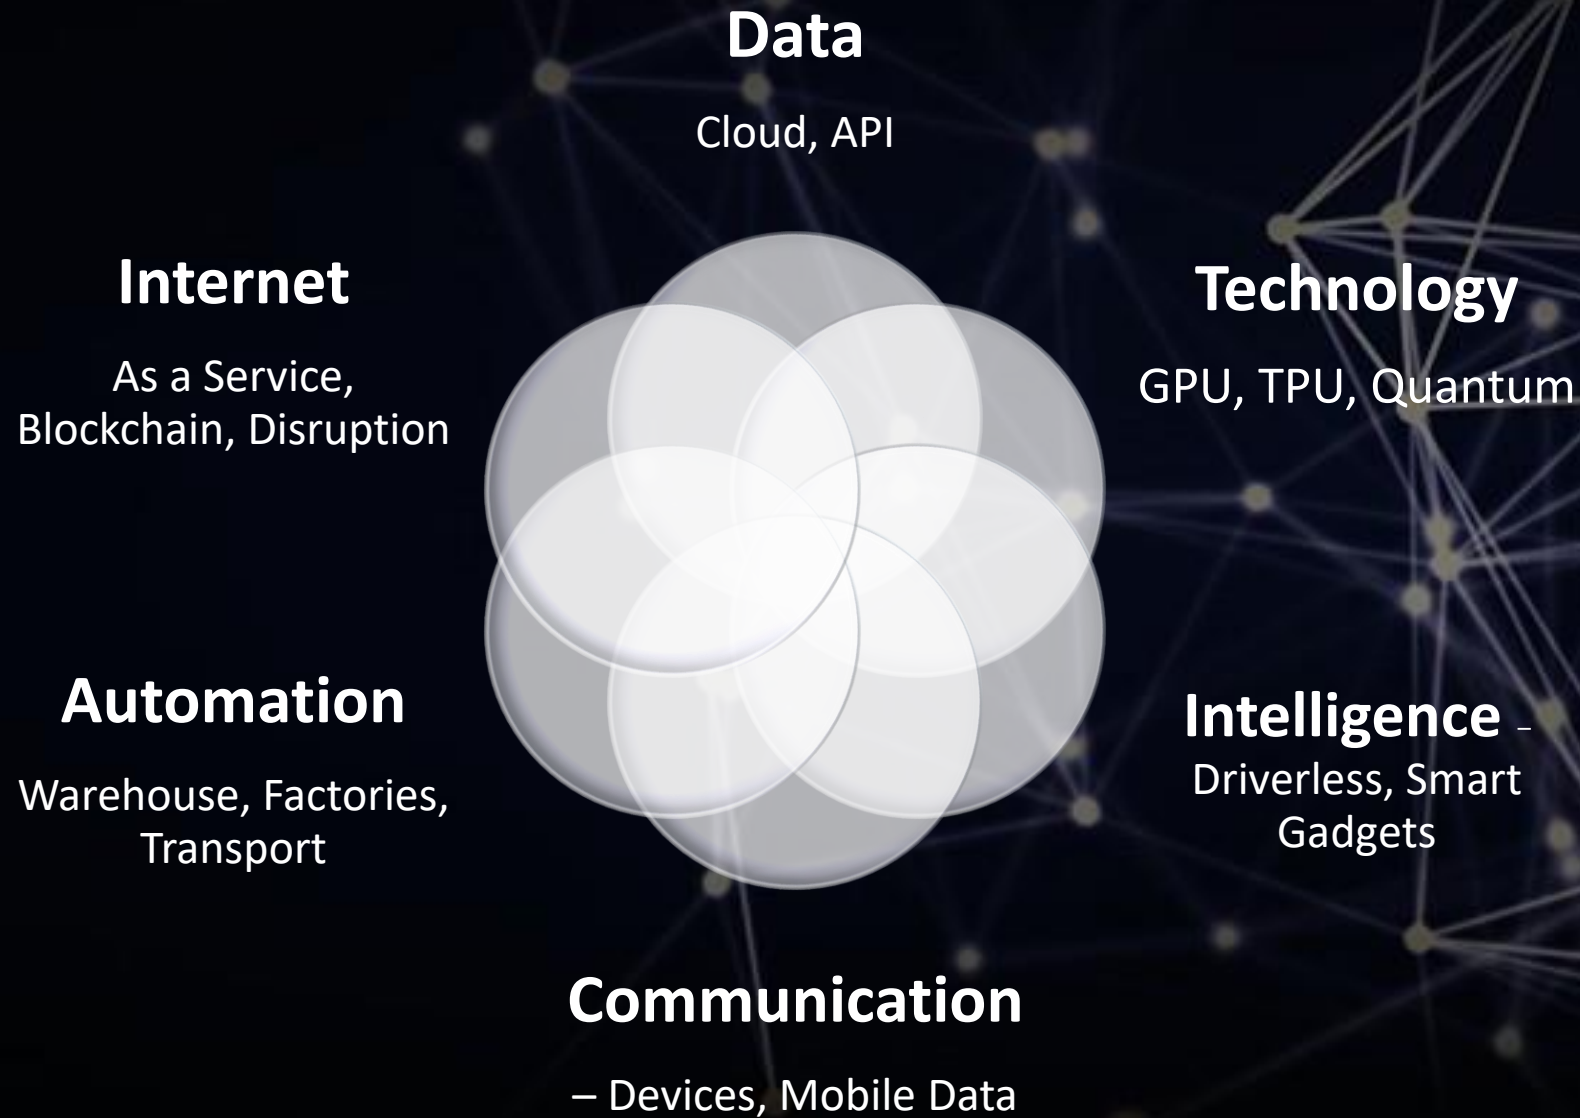
Artificial Intelligence – Computers with the ability to reason as humans

Machine Learning – Computers with the ability to learn without being explicitly programmed


Deep Learning – Network capable of adapting itself to new data

- 1956, the term “AI” coined by scientist John McCarthy at Dartmouth College
- Around 1980 ML starts to flourish
 - IBM’s Deep Blue computer beats Russian chess master Garry Kasparov in 1997
 - IBM’s Watson beat human players on Jeopardy in 2011 utilizing ML
- 2010 – Deep Learning flourishes
 - 2017 – Densenet is announced
- Today everyone is chasing AI as the new Gold-Rush


Convergence




Reality vs Expectation

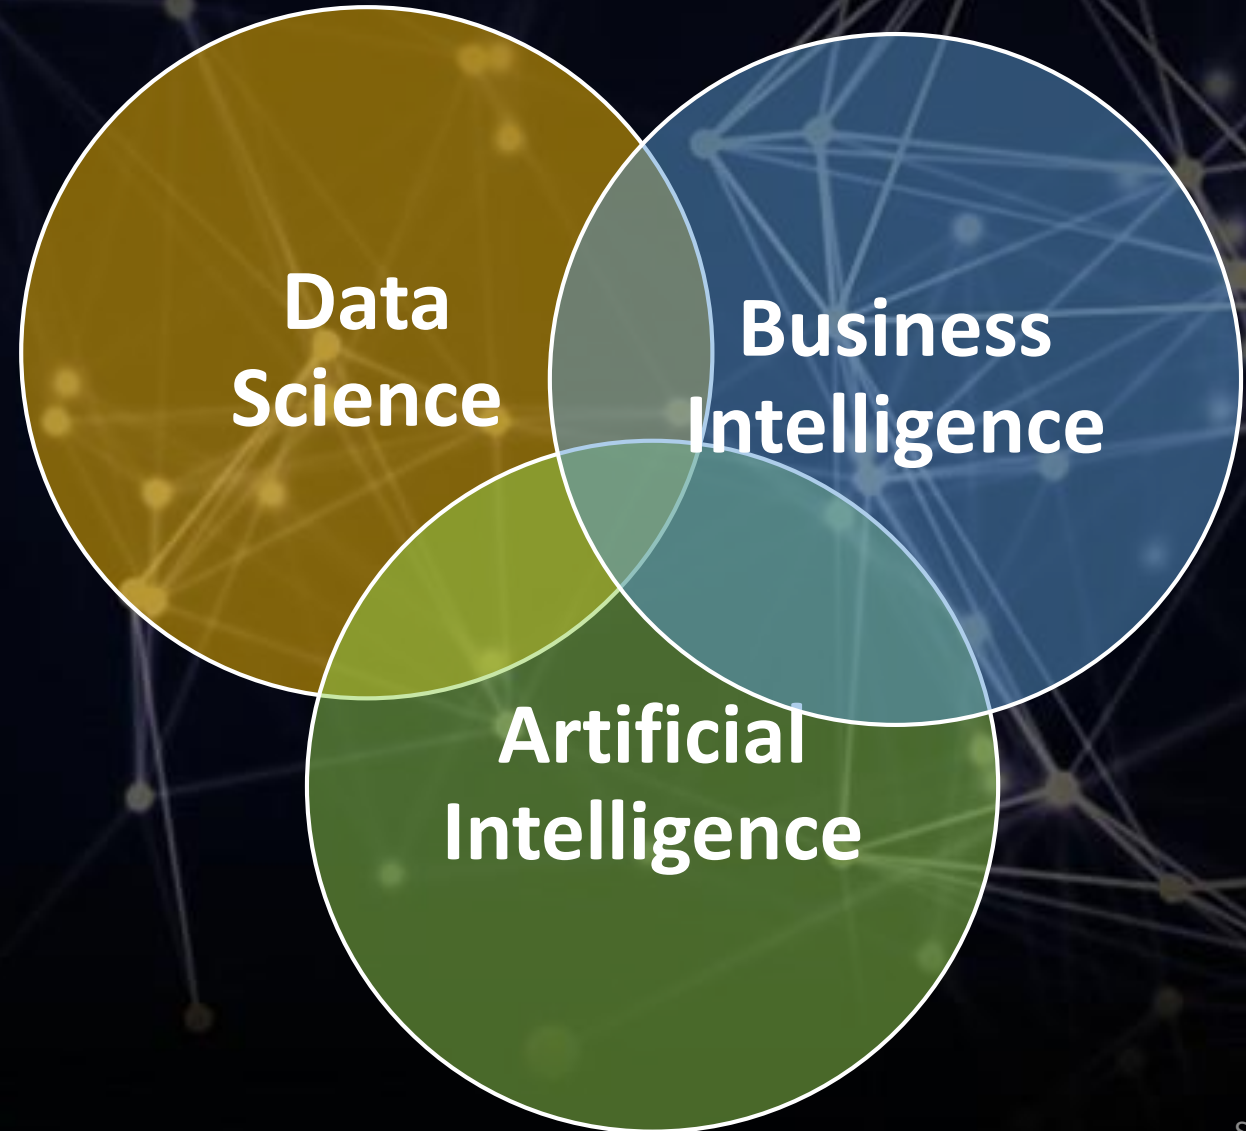
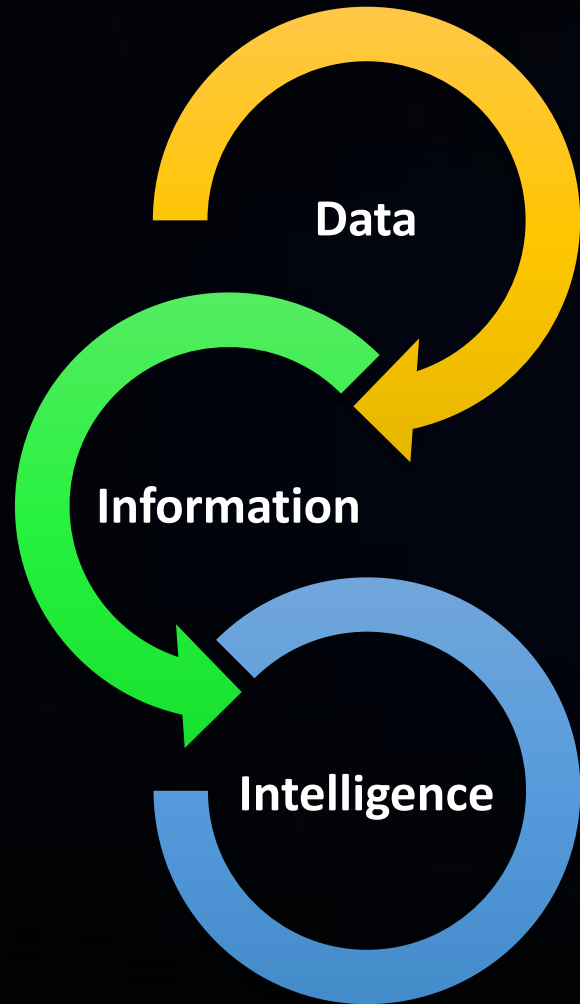


Don't worry, we are okay for now!





Talk is cheap, show me your intelligence!



Experiment – how do we learn?



Machine Learning in ML Studio

Anomaly Detection

- One-class Support Vector Machine
- Principal Component Analysis-based Anomaly Detection
- Time Series Anomaly Detection*

Classification

Two-class Classification

- Averaged Perceptron
- Bayes Point Machine
- Boosted Decision Tree
- Decision Forest
- Decision Jungle
- Logistic Regression
- Neural Network
- Support Vector Machine

Multi-class Classification

- Decision Forest
- Decision Jungle
- Logistic Regression
- Neural Network
- One-vs-all

Clustering

- K-means Clustering

Recommendation

- Matchbox Recommender

Regression

- Bayesian Linear Regression
- Boosted Decision Tree
- Decision Forest
- Fast Forest Quantile Regression
- Linear Regression
- Neural Network Regression
- Ordinal Regression
- Poisson Regression

Statistical Functions

- Descriptive Statistics
- Hypothesis Testing T-Test
- Linear Correlation
- Probability Function Evaluation

Text Analytics

- Feature Hashing
- Named Entity Recognition
- Vowpal Wabbit

Computer Vision

- OpenCV Library

Guest Access Workspace: Free trial access without logging in.
 Free Workspace: Free persisted access, no Azure subscription needed.
 Standard Workspace: Full access with SLA under an Azure subscription.

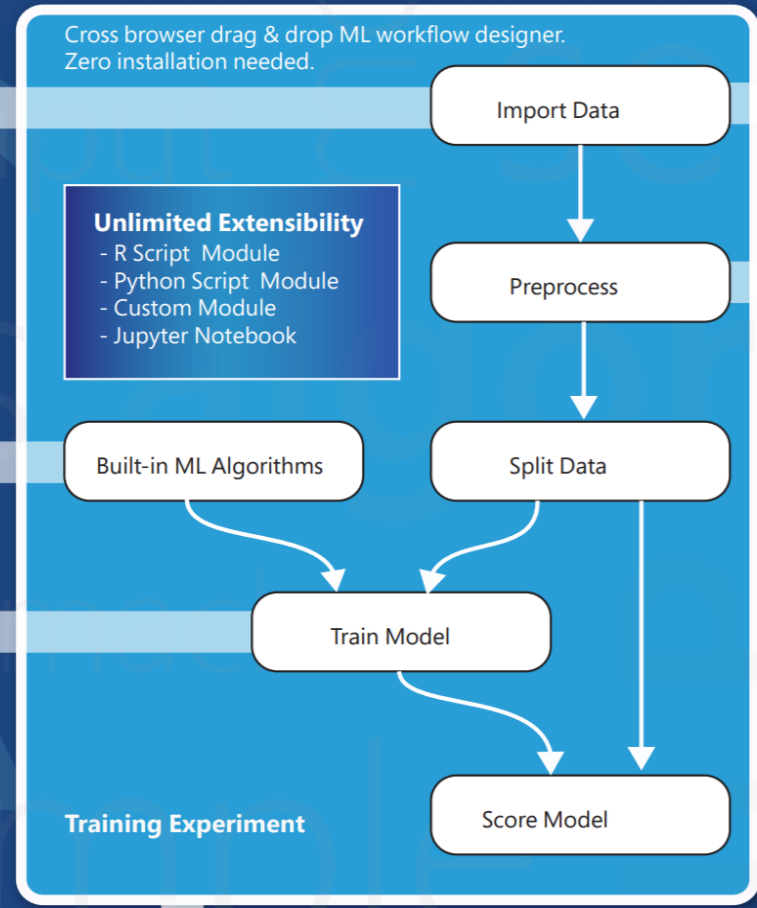
Data Source	Data Format
- Azure Blob Storage	- ARFF
- Azure SQL DB	- CSV
- Azure SQL DW*	- SVMLight
- Azure Table	- TSV
- Desktop Direct Upload	- Excel
- Hadoop Hive Query	- ZIP
- Manual Data Entry	
- OData Feed	
- On-prem SQL Server*	
- Web URL (HTTP)	

Data/Model Visualization

- Scatterplots
- Bar Charts
- Box plots
- Histogram
- R and Python Plotting Libraries
- REPL with Jupyter Notebook
- ROC, Precision/Recall, Lift
- Confusion Matrix
- Decision Tree*

Training

- Cross Validation
- Retraining
- Parameter Sweep



Data Preparation

- Clean Missing Data
- Clip Outliers
- Edit Metadata
- Feature Selection
- Filter
- Learning with Counts
- Normalize Data
- Partition and Sample
- Principal Component Analysis
- Quantize Data
- SQLite Transformation
- Synthetic Minority Oversampling Technique

Enterprise Grade Cloud Service

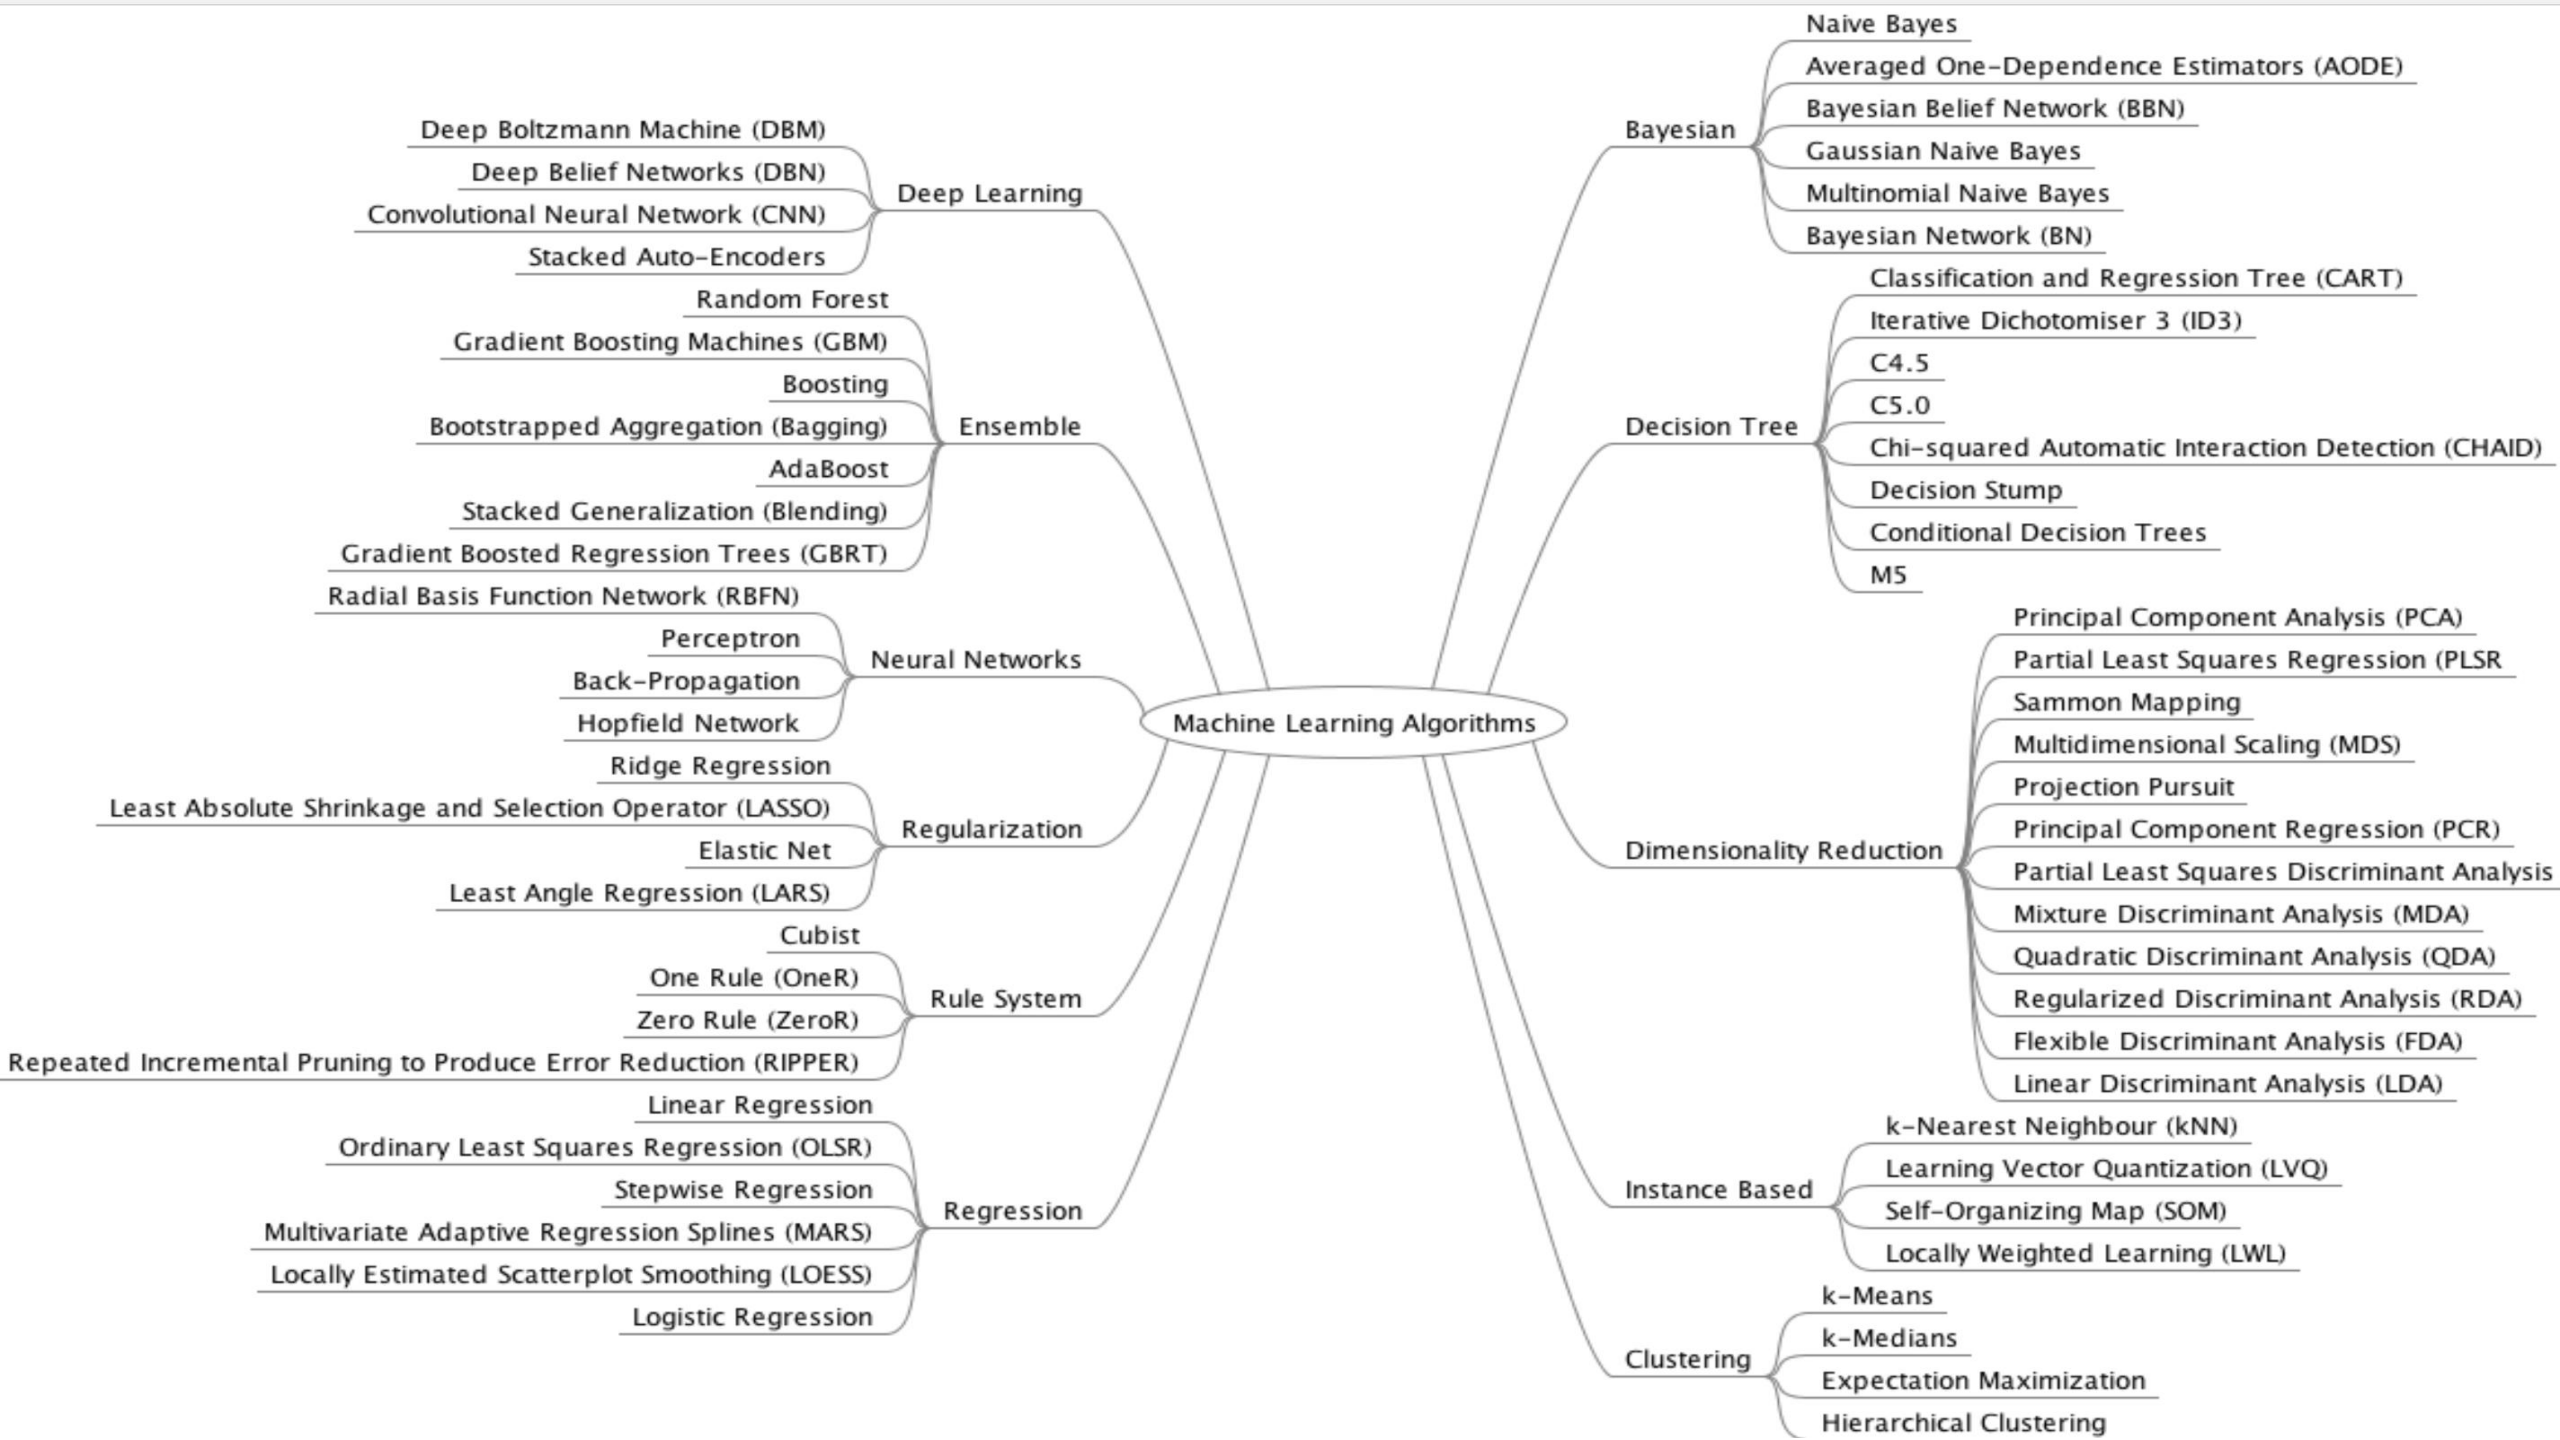
- SLA: 99.95% Guaranteed Up-time
- Azure AD Authentication
- Compute at Large Scale
- Multi-geo Availability
- Regulatory Compliance*

One-click Operationalization



Community

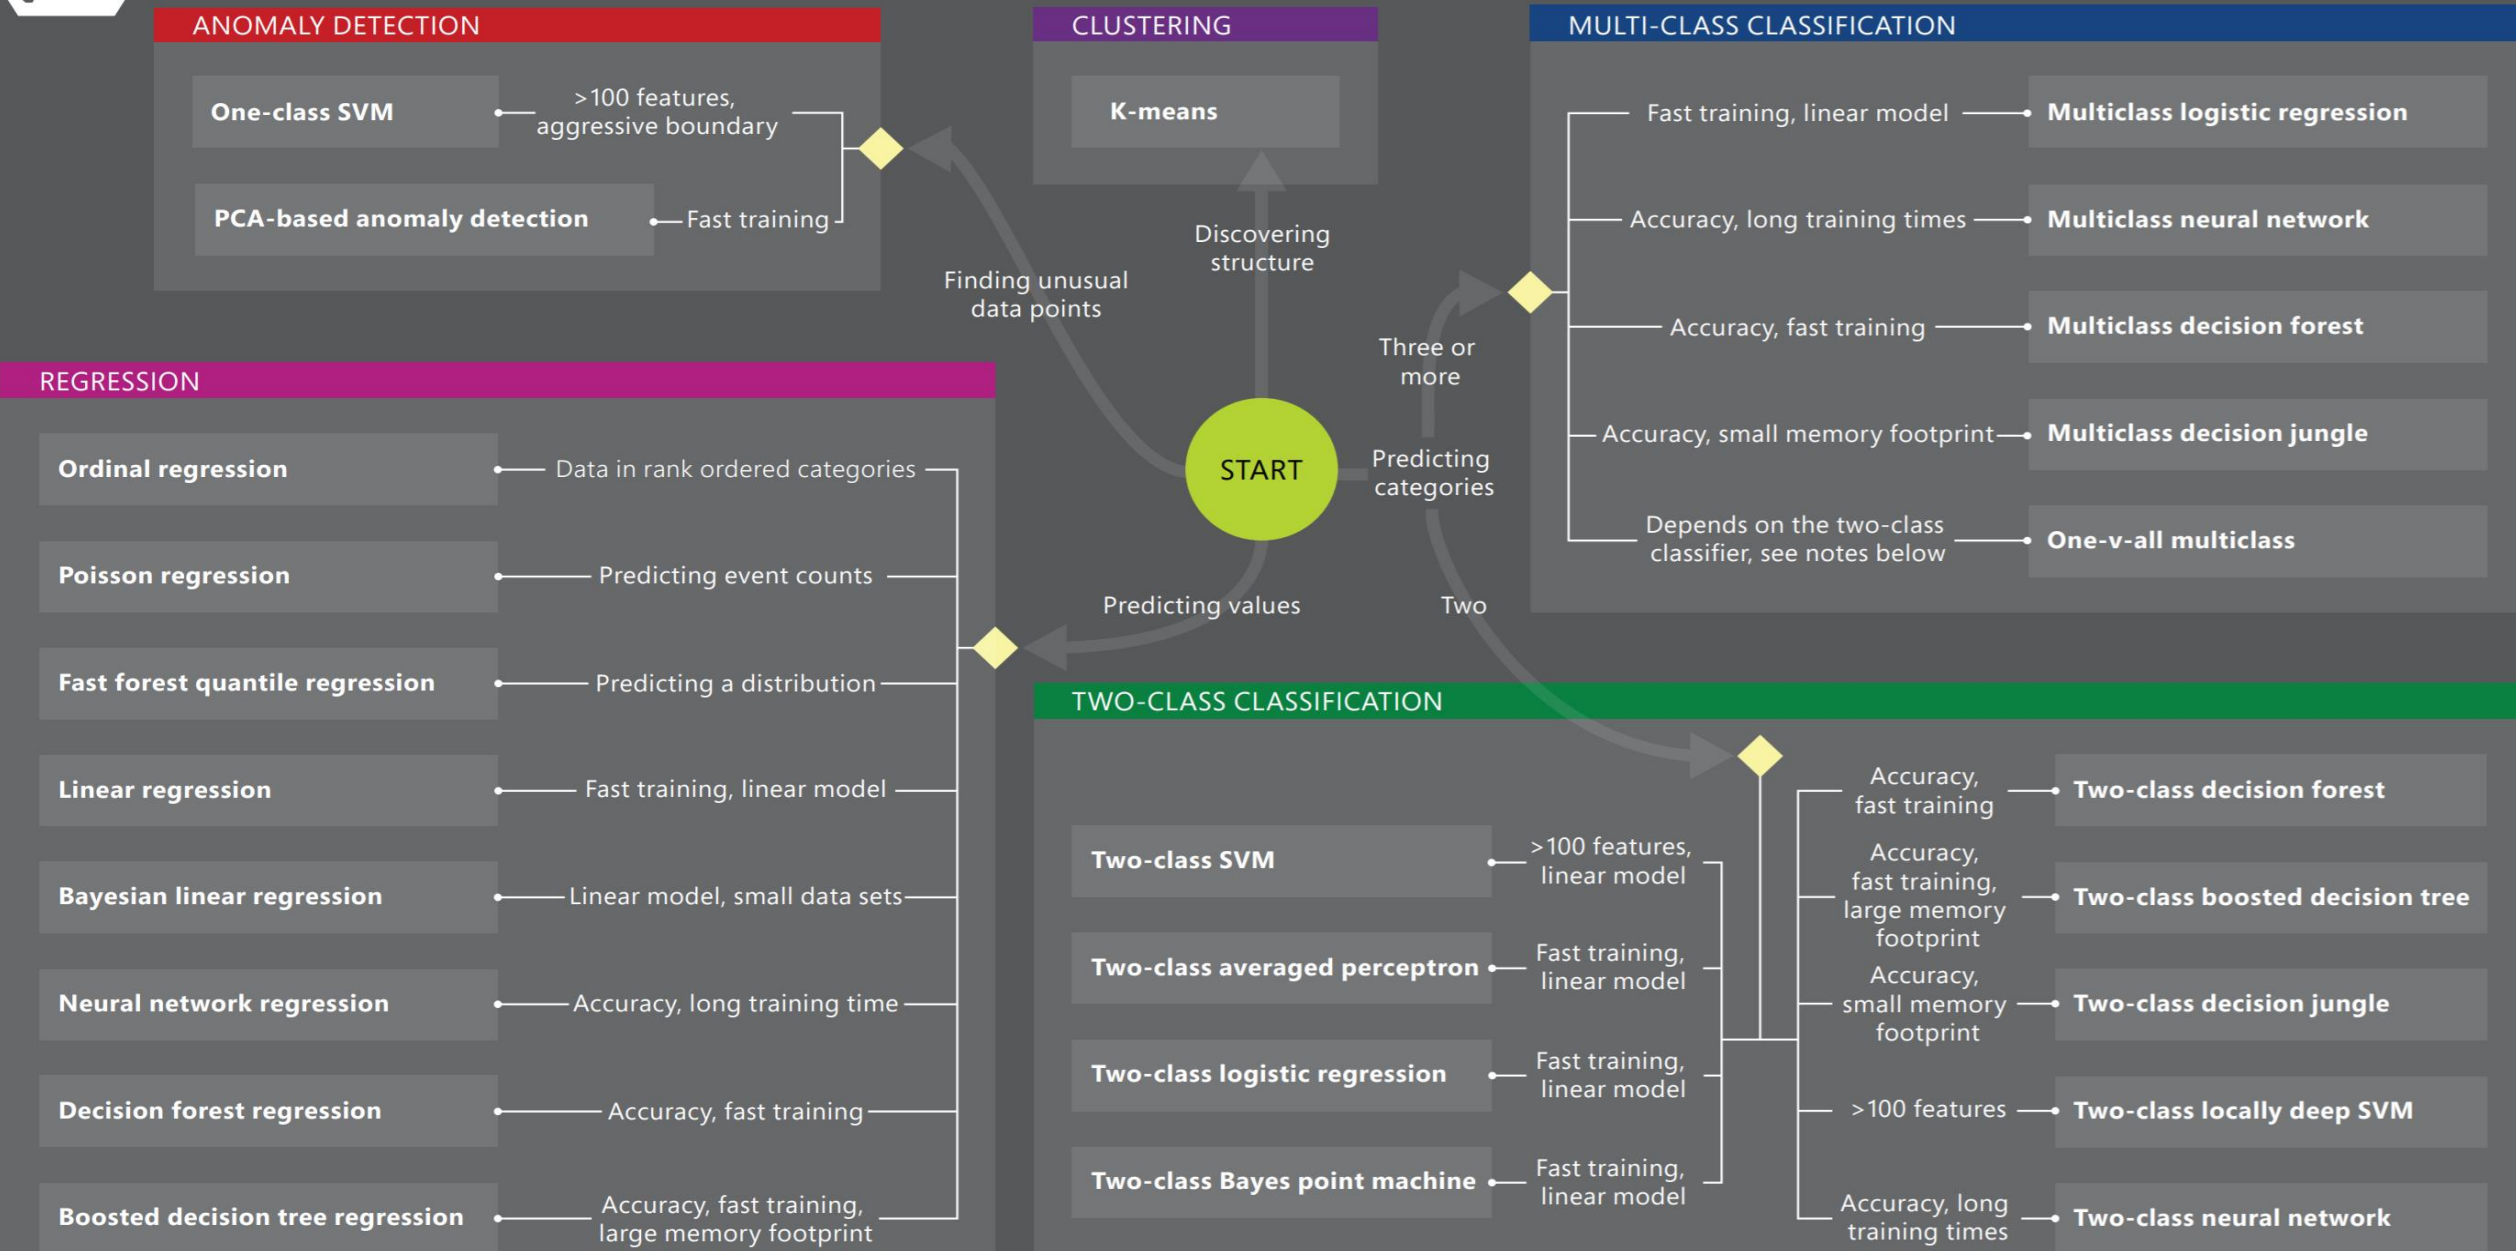
- Gallery (<http://gallery.azureml.net>)
- Samples & Templates
- Workspace Sharing and Collaboration
- Live Chat & MSDN Forum Support





Microsoft Azure Machine Learning: Algorithm Cheat Sheet

This cheat sheet helps you choose the best Azure Machine Learning Studio algorithm for your predictive analytics solution. Your decision is driven by both the nature of your data and the question you're trying to answer.





- <https://docs.microsoft.com/en-us/azure/machine-learning/service/overview-what-is-azure-ml>
- <https://docs.microsoft.com/en-us/azure/machine-learning/desktop-workbench/tutorial-classifying-iris-part-1>
- <https://www.kdnuggets.com/2016/10/clustering-key-terms-explained.html>
- <https://www.kaggle.com/kanncaa1/machine-learning-tutorial-for-beginners>
- <https://github.com/rasbt/python-machine-learning-book-2nd-edition>
- <https://github.com/eriklindernoren/ML-From-Scratch>
- <https://www.altexsoft.com/blog/datascience/comparing-machine-learning-as-a-service-amazon-microsoft-azure-google-cloud-ai/>

Thank you!



Contact:

www.Ai-Tech.uk

AI Tech UK

@AITechUG #AIUserGroup

Mathewmatics@outlook.com

@smdisrupt

