Demystifying Machine Learning

Tom Gielow
Vice President Technology & Architecture

Google Trends Data

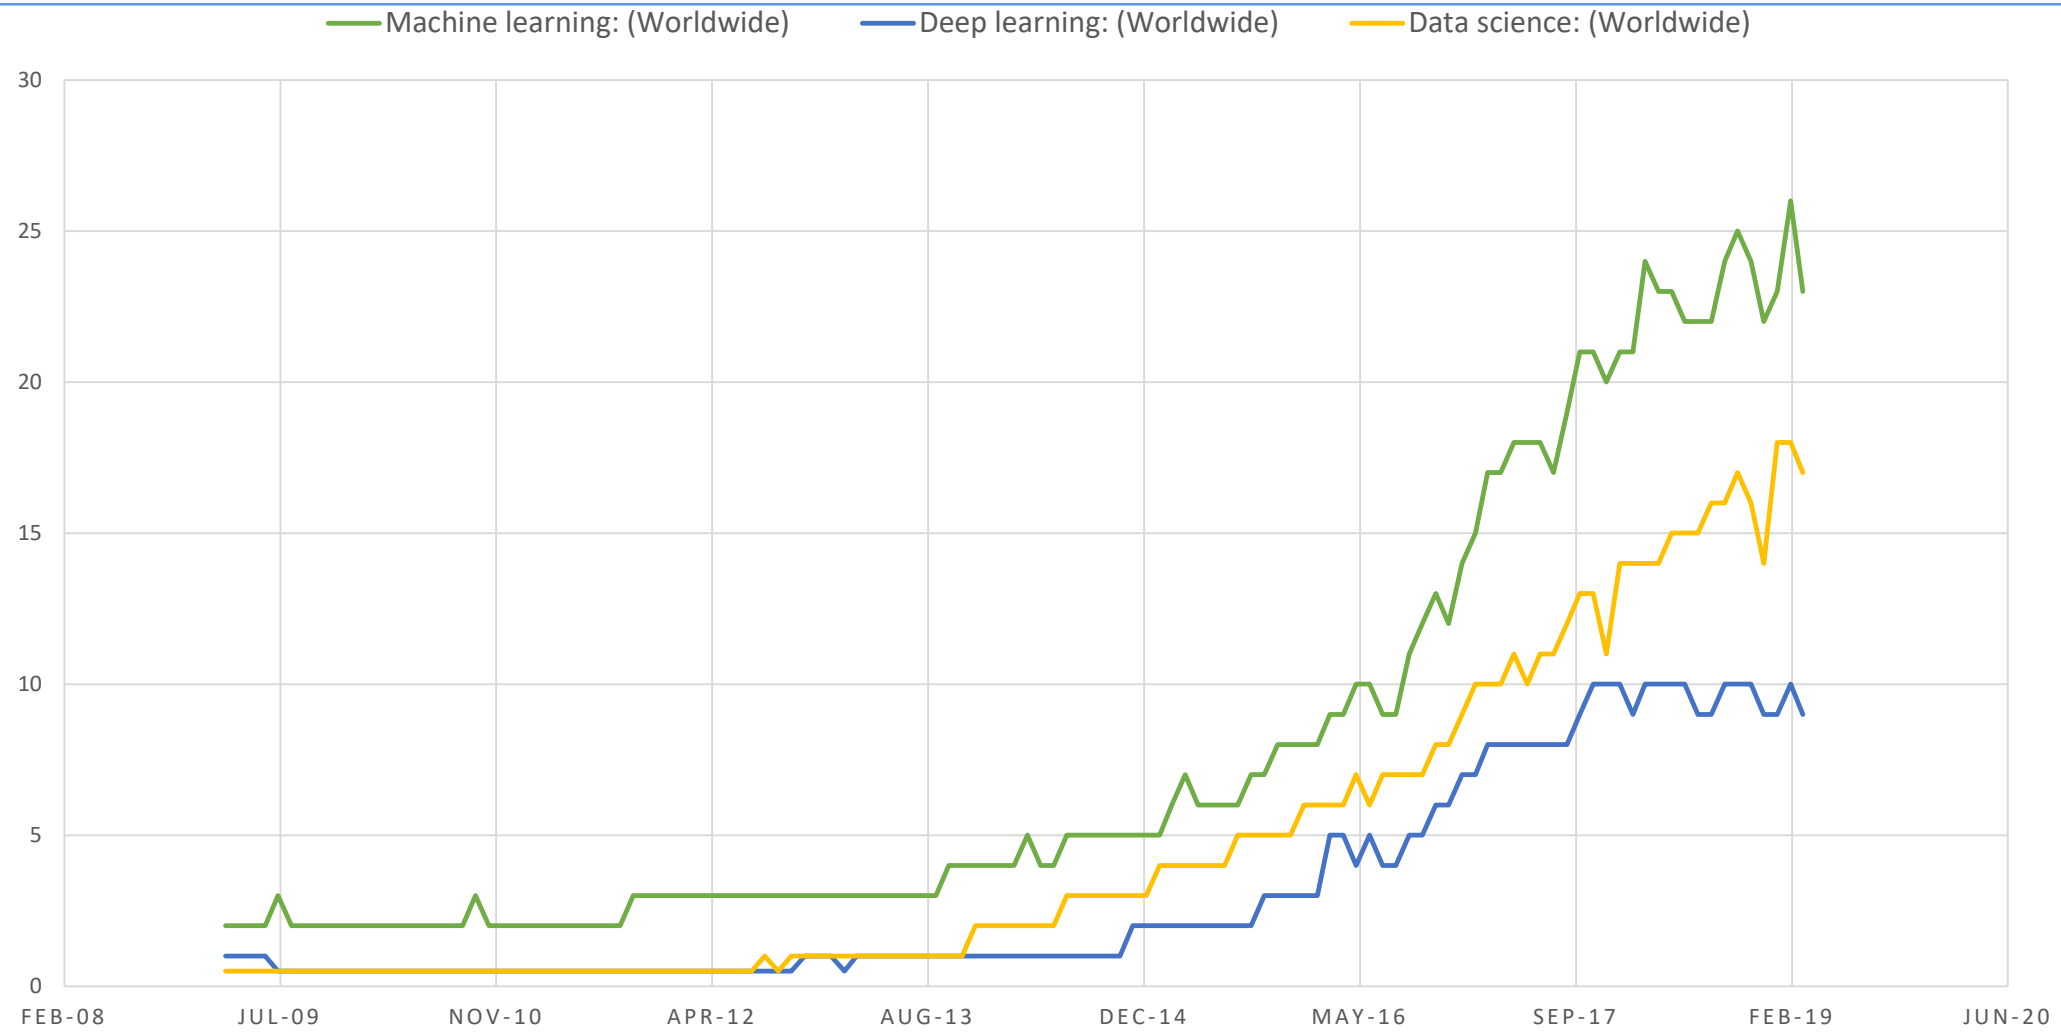

Google Trends Data

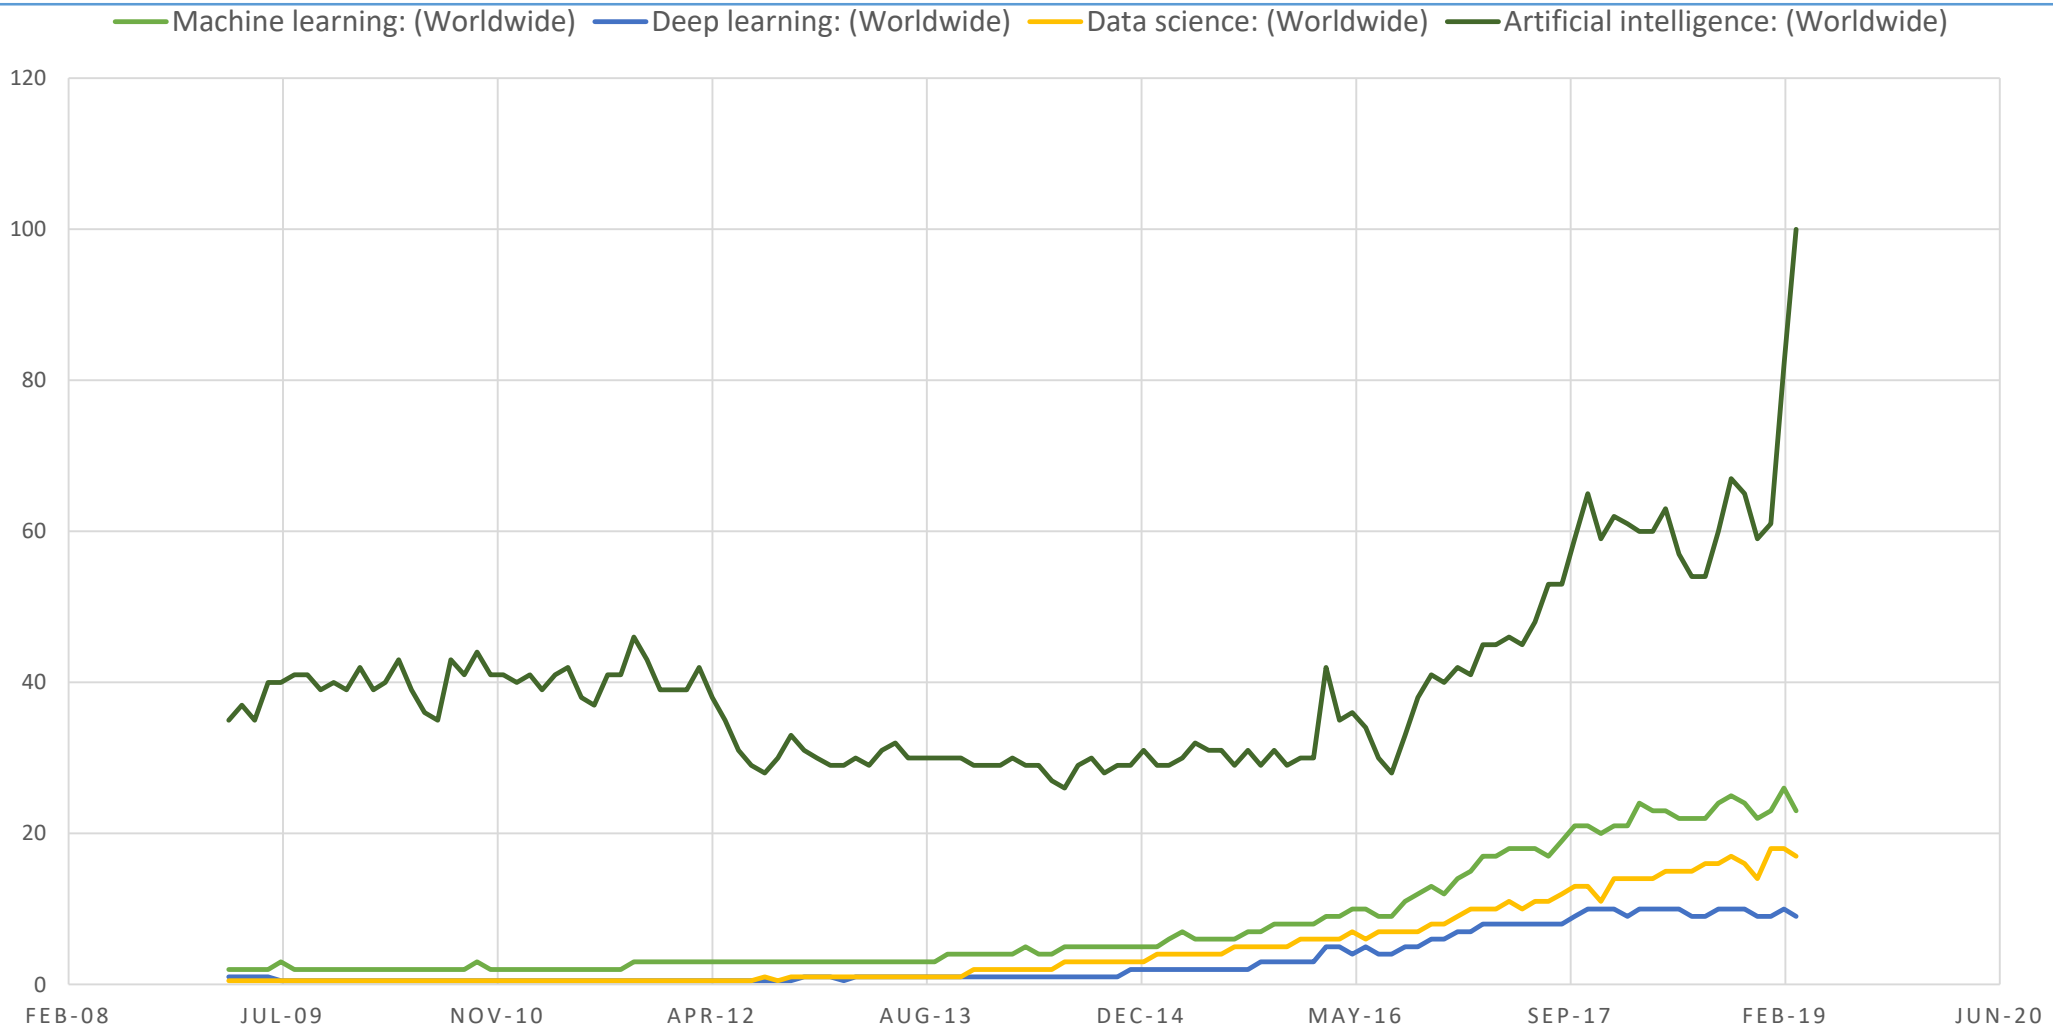

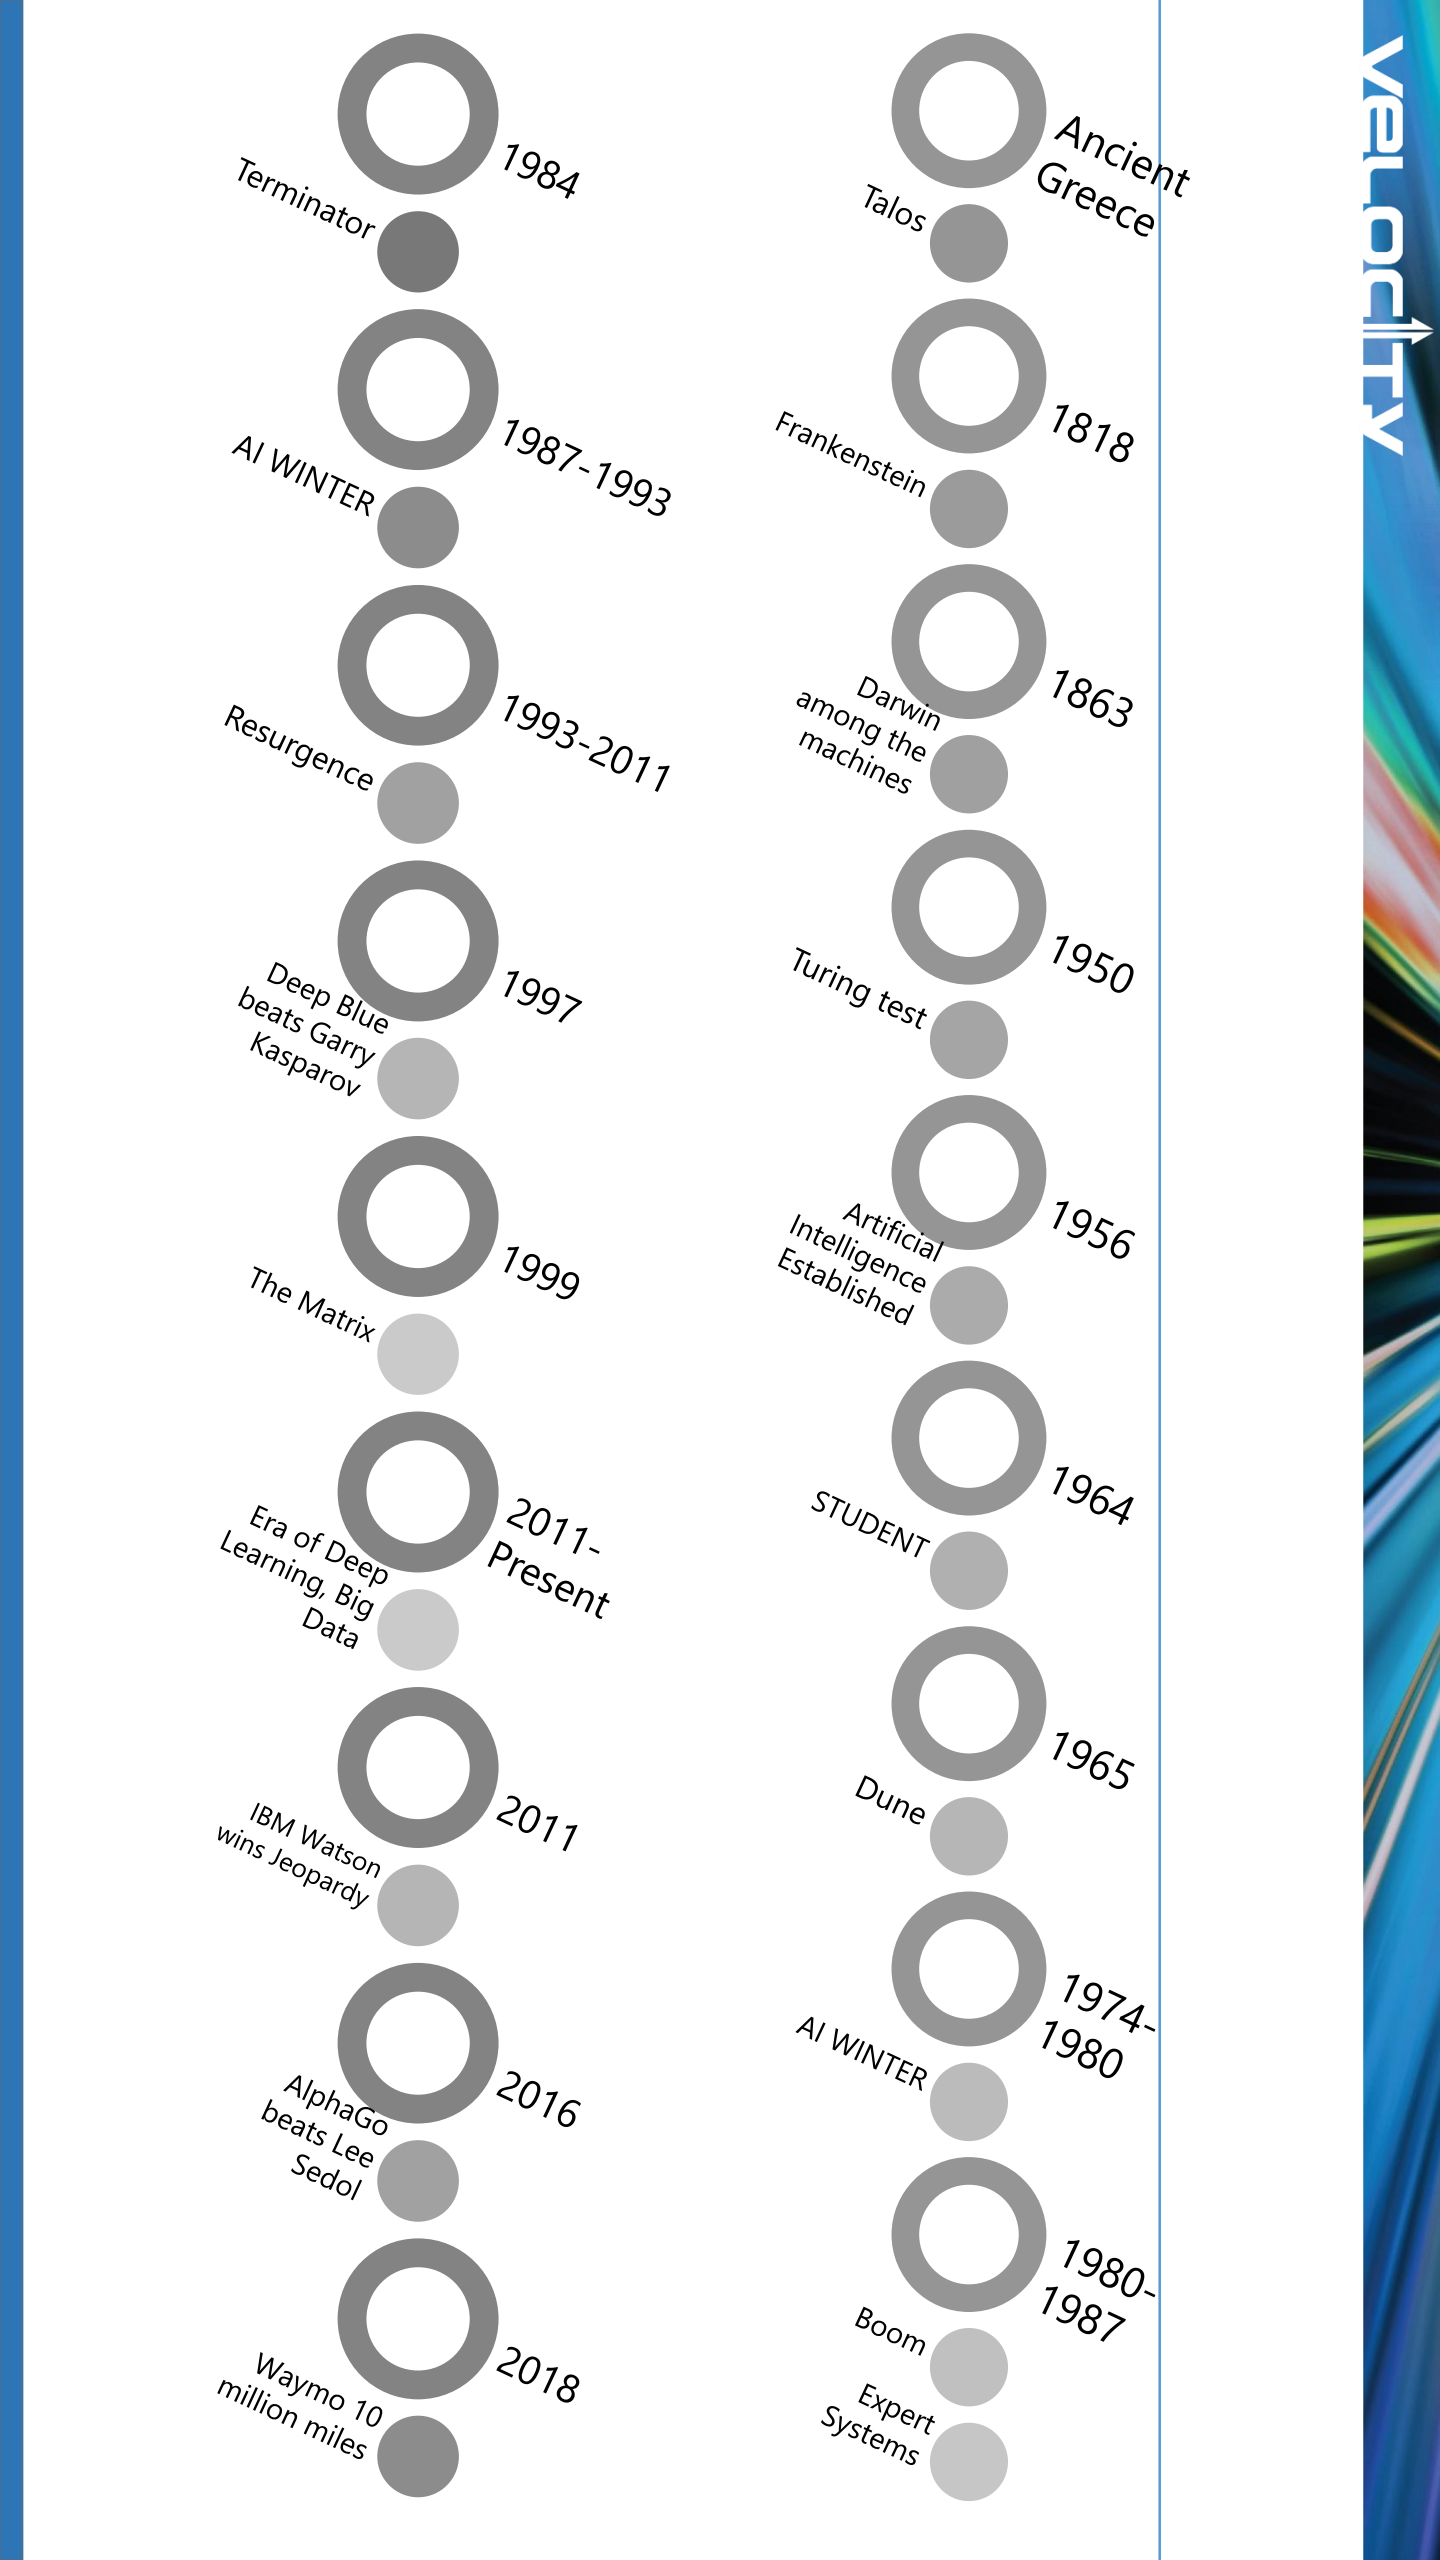
Ancient Greece
Talos
1818
Frankenstein
1863
Darwin among the machines
1950
Turing test
1956
Artificial Intelligence Established
1964
STUDENT
1965
Dune
1974-1980
AI WINTER
1980-1987
Boom
Expert Systems
1984
Terminator
1987-1993
AI WINTER
1993-2011
Resurgence
1997
Deep Blue beats Garry Kasparov
1999
The Matrix
2011-Present
Era of Deep Learning, Big Data
2011
IBM Watson wins Jeopardy
2016
AlphaGo beats Lee Sedol
2018
Waymo 10 million miles

What is Artificial Intelligence?

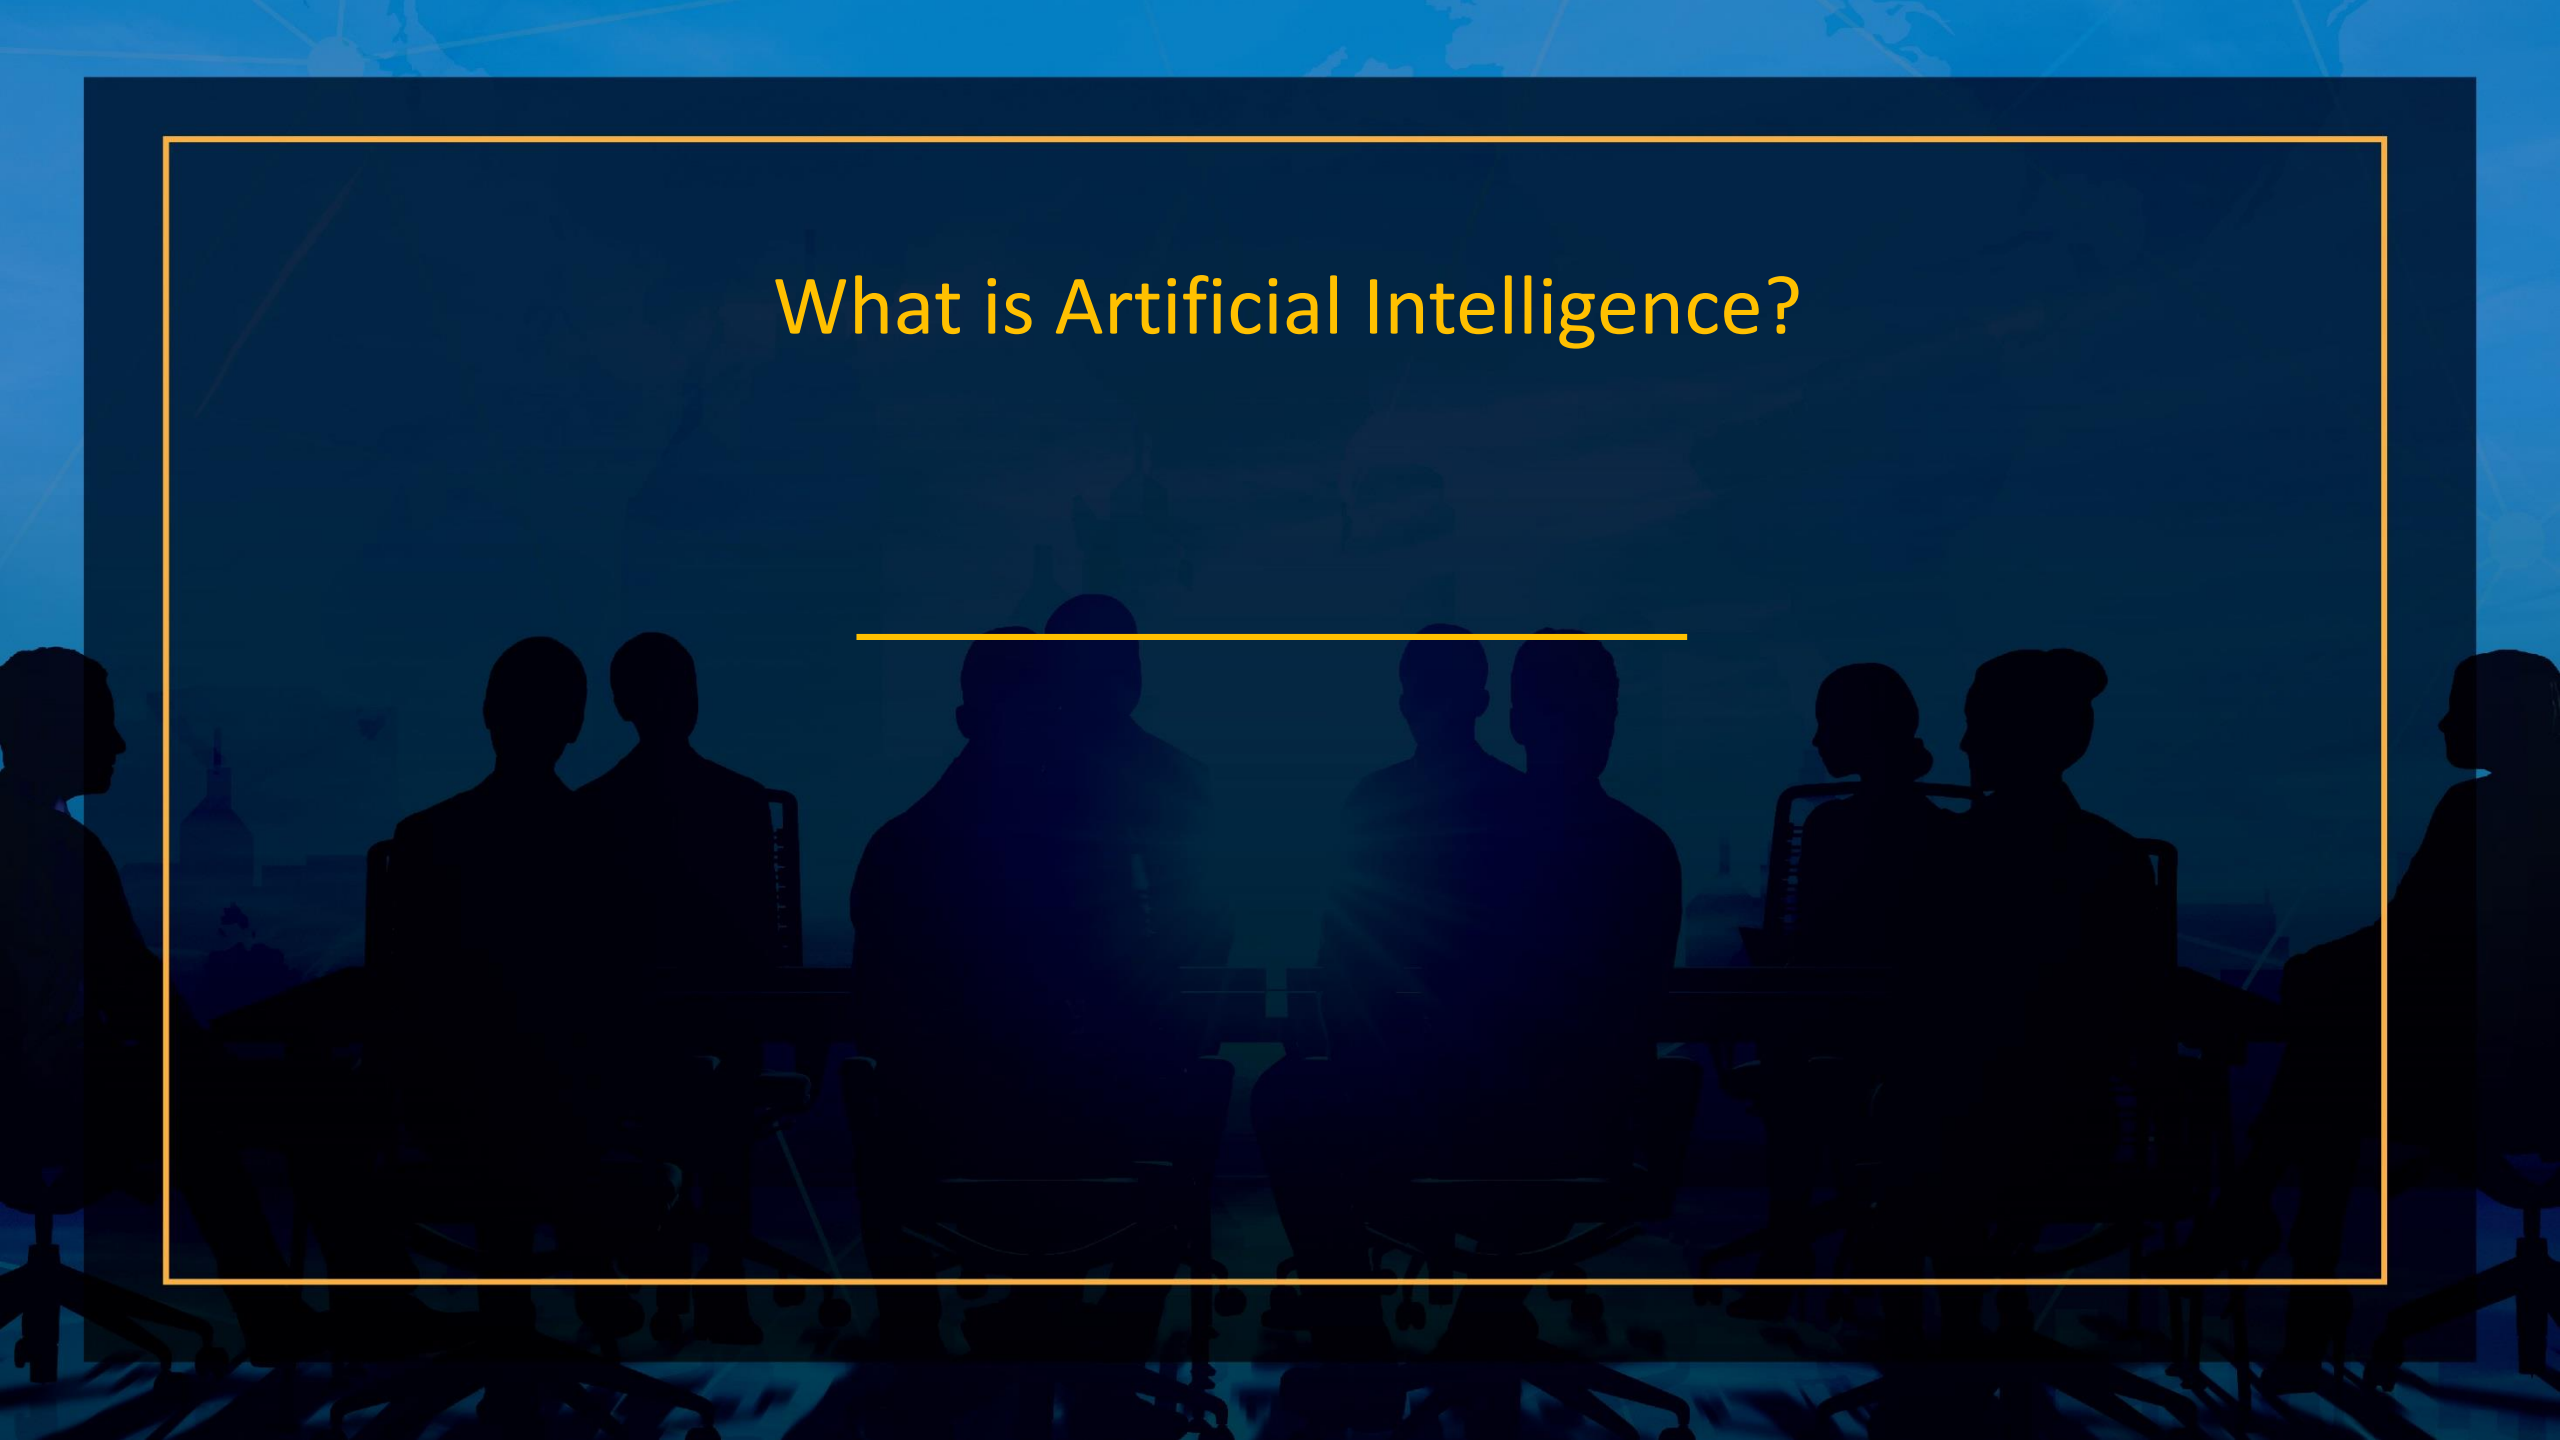

What is Artificial Intelligence?

- **Artificial Intelligence:** The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.

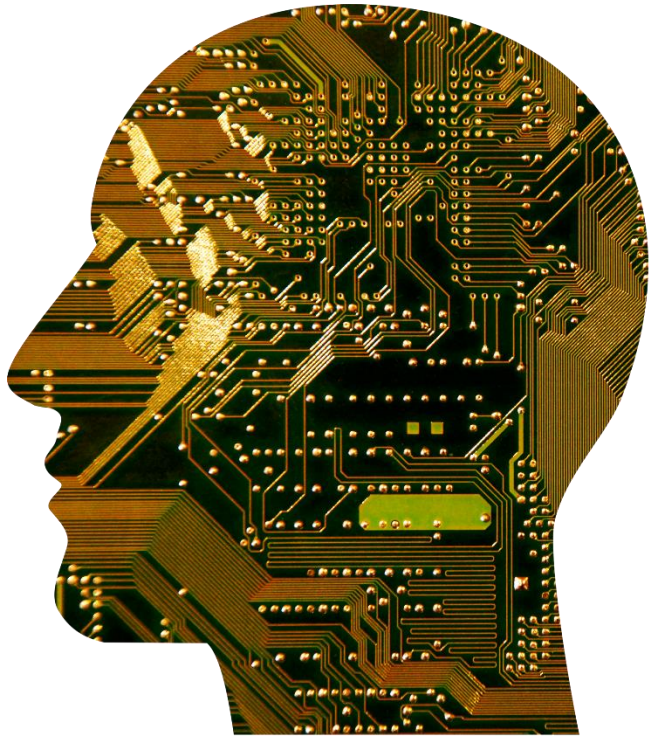

Strong vs. Weak Artificial Intelligence

STRONG AI

- Human-like broad intelligence
- Understands problems and context
- Examples
  - HAL 9000
  - C3PO
  - Lt. Commander Data

WEAK AI

- Limited to specific tasks
- Solve complex problems without understanding them
- Examples
  - Match search terms to pages
  - Recommend TV shows based on views

Machine Learning

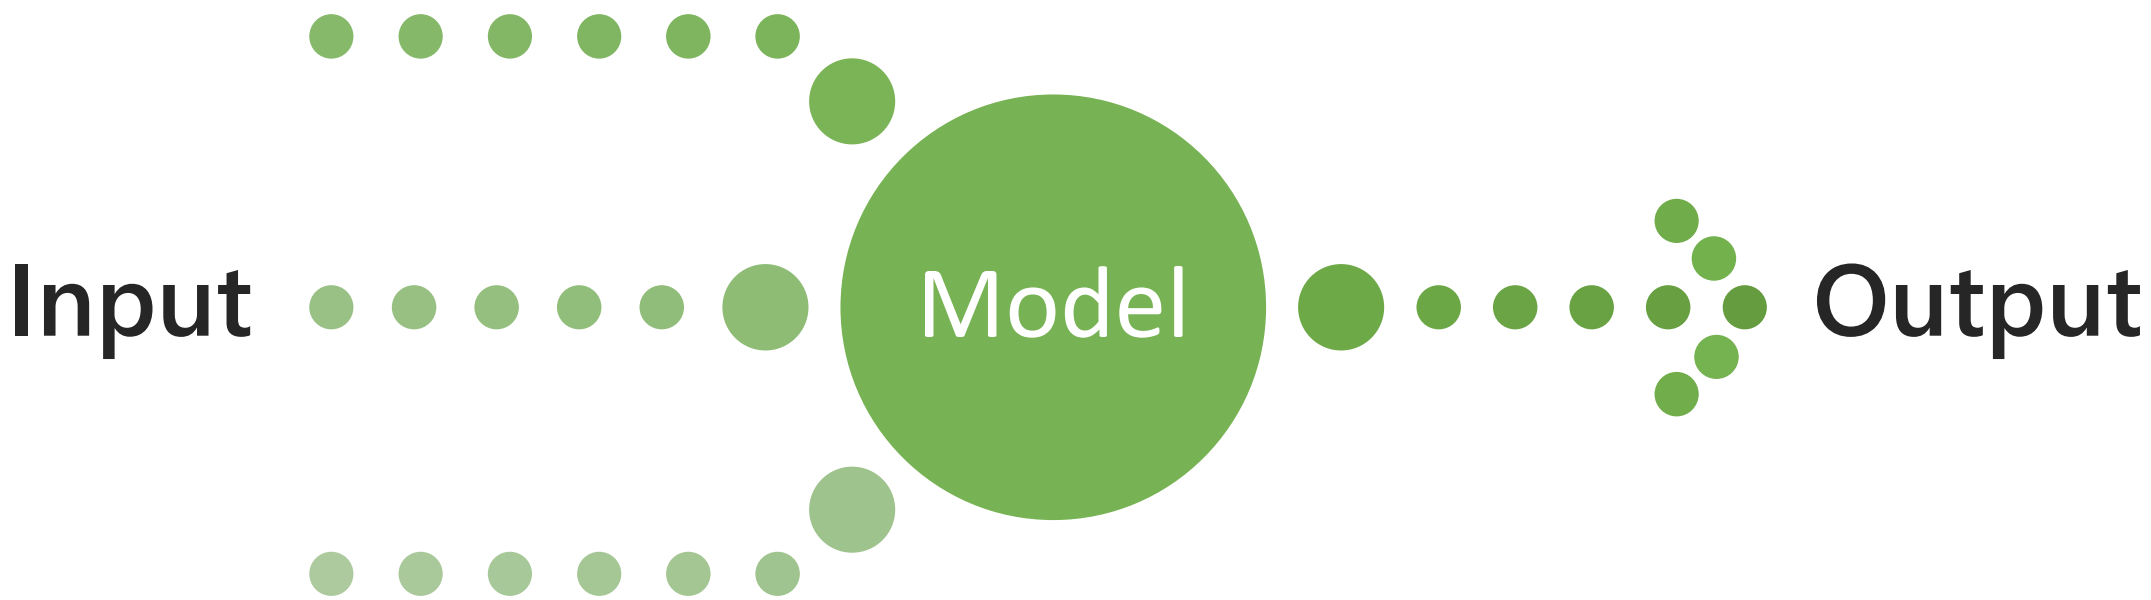

Building a Machine Learning Model

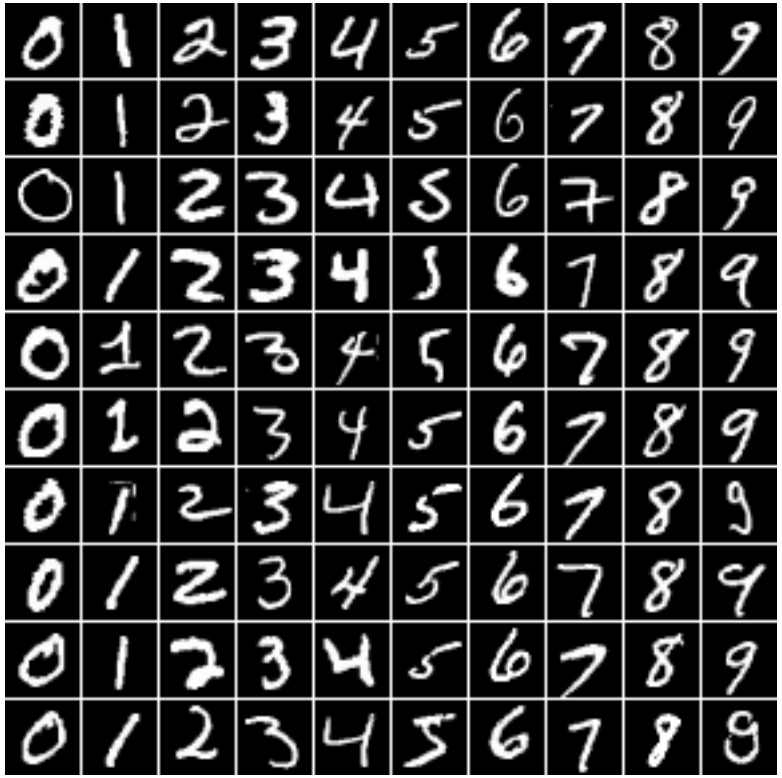

- MNIST dataset
  - 70 000 handwritten digits

Building a Machine Learning Model

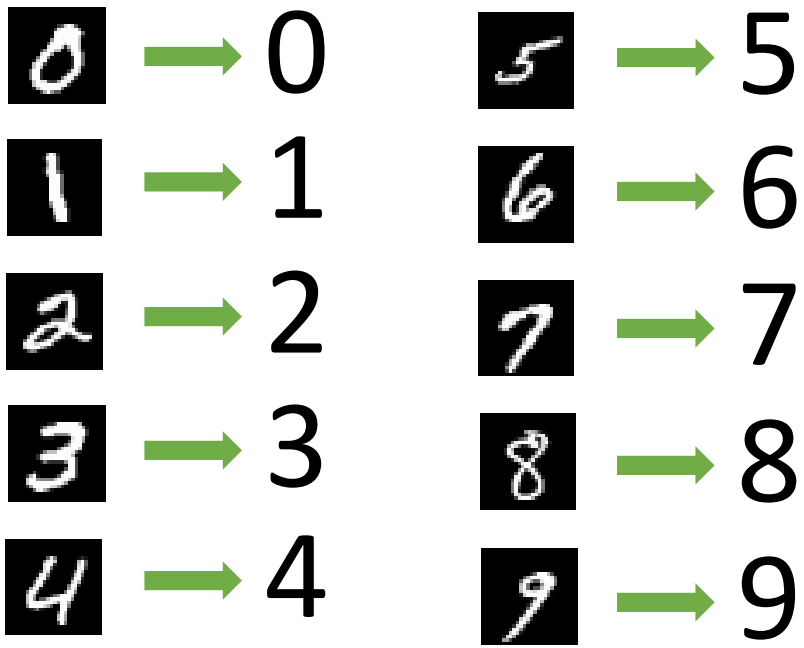

- MNIST dataset
  - 70 000 handwritten digits
- Each image is labeled according to its contents

Building a Machine Learning Model

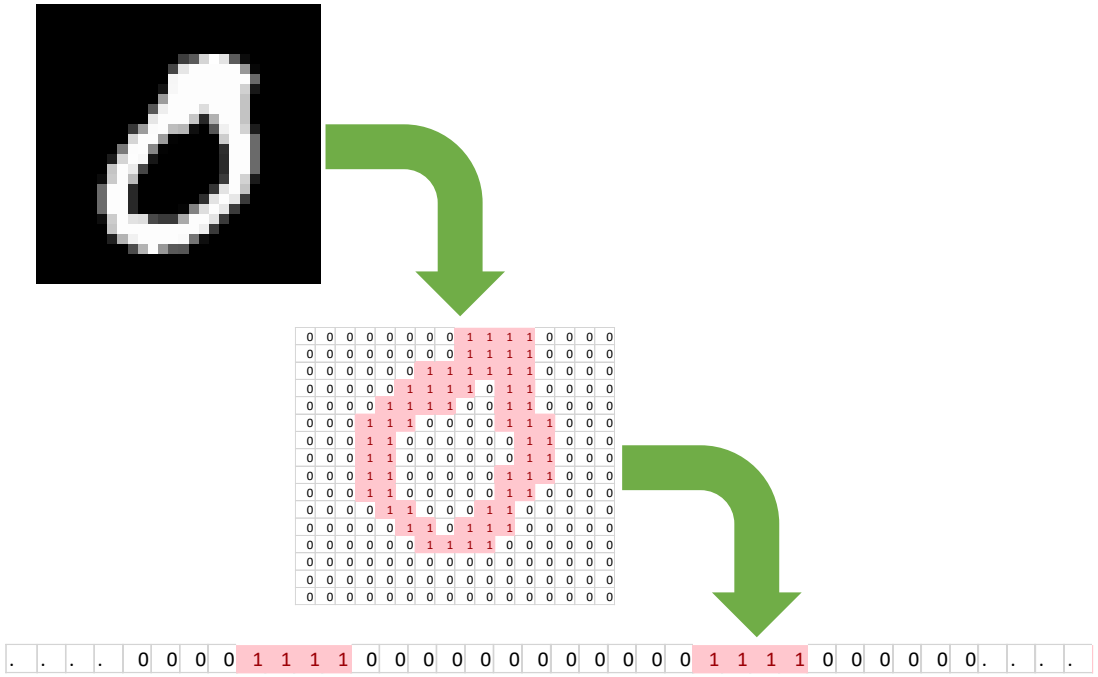

- MNIST dataset
  - 70 000 handwritten digits
- Each image is labeled according to its contents
- Data that represents images with label fed to ML algorithm

Building a Machine Learning Model

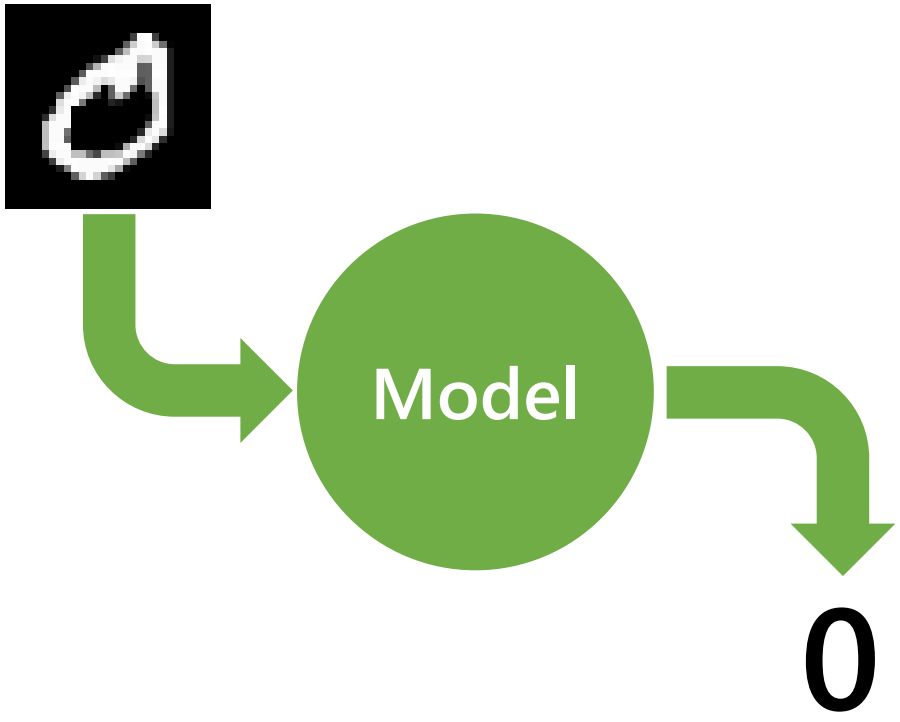

- MNIST dataset
  - 70 000 handwritten digits
- Each image is labeled according to its contents
- Data that represents images with label fed to ML algorithm
- Model is produced that can recognize handwritten digits

Unsupervised Learning: Clustering

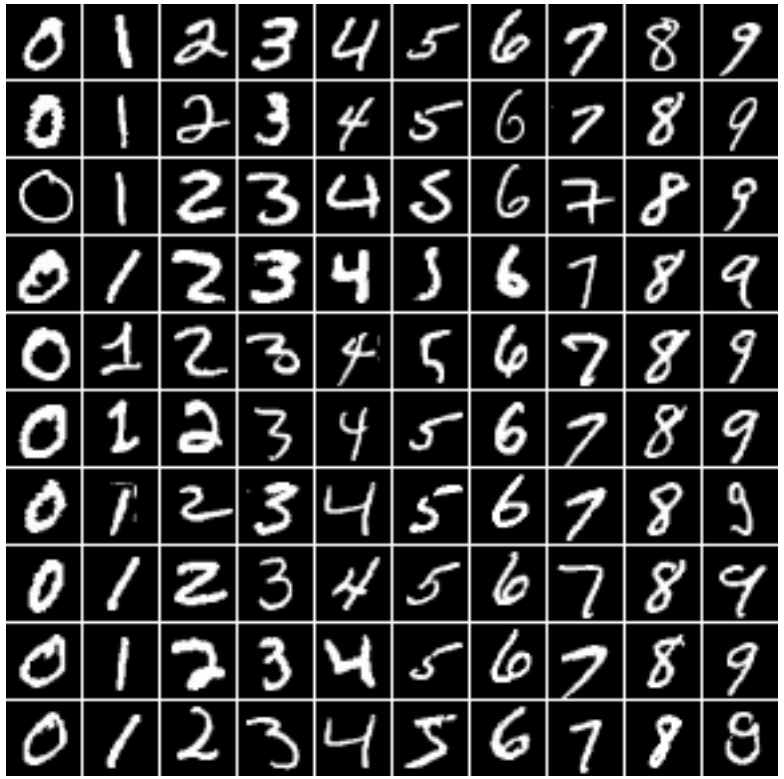

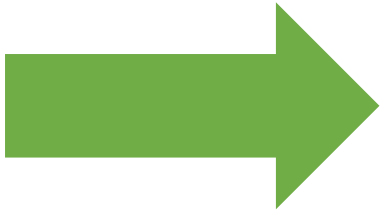

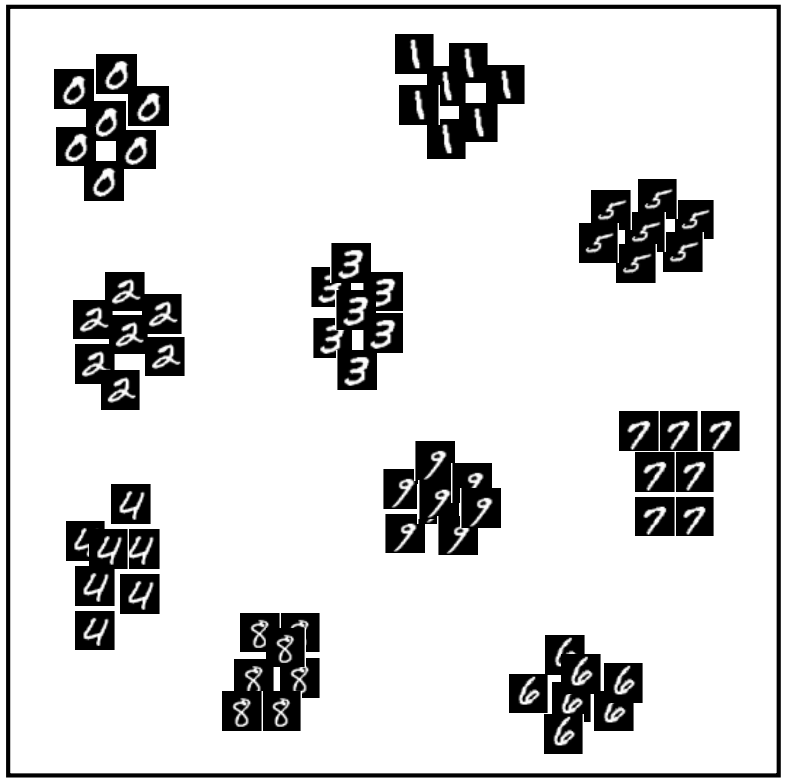

Types of Machine Learning Problems

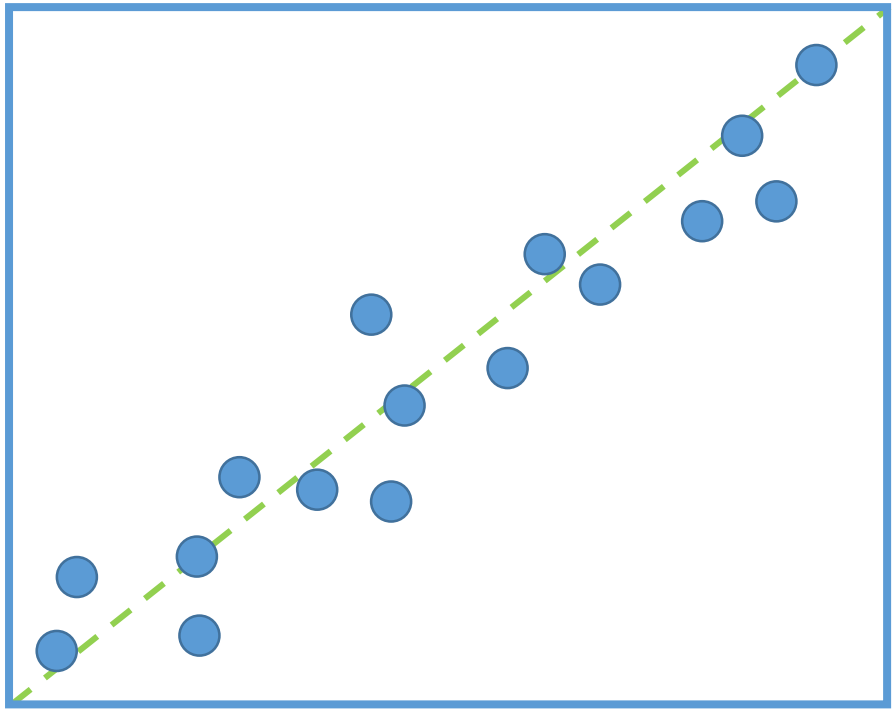

Regression

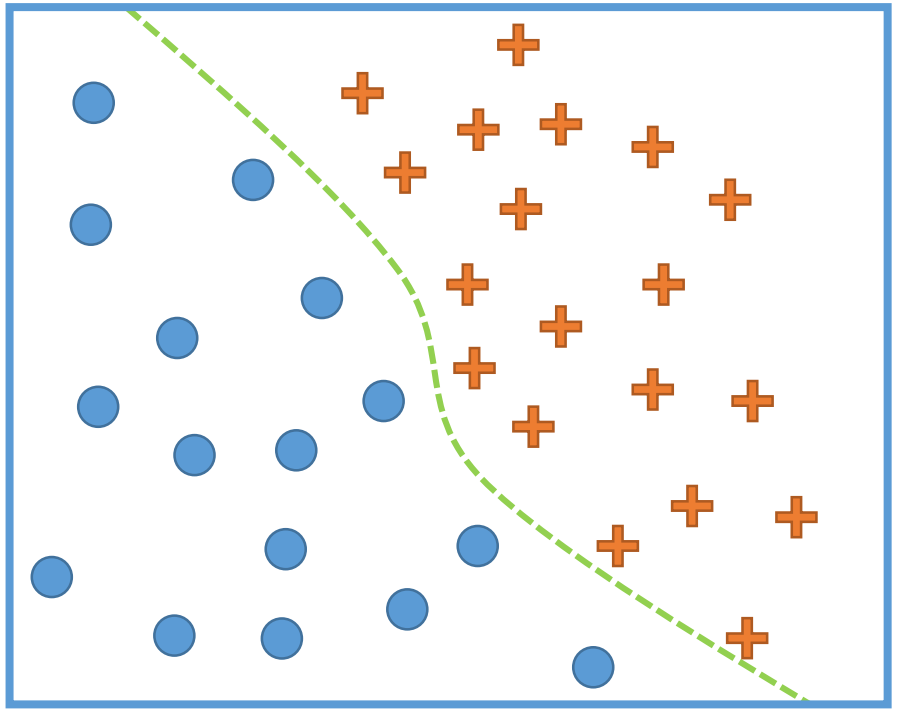

Classification

Deep Learning

- Usually an artificial neural network
- Extracts important information from data automatically
- Potential for better results than other ML methods
- Require large amounts of data
- Computationally intensive

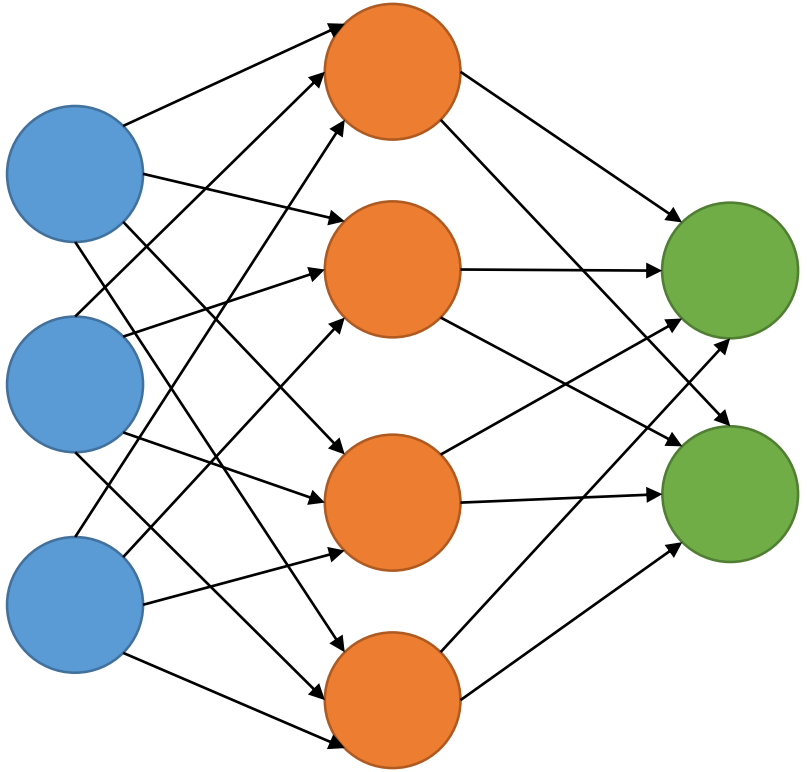

Artificial Intelligence Hierarchy

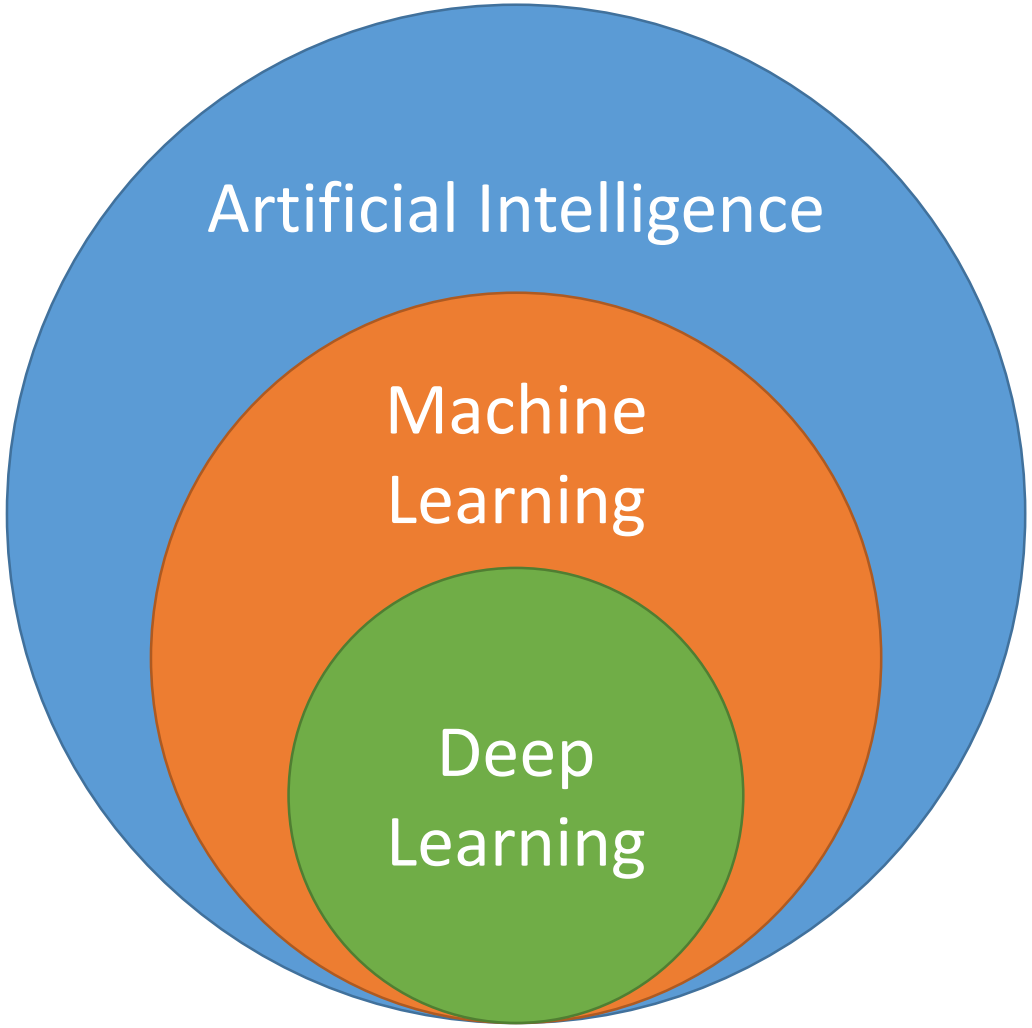

Machine Learning Examples

- Personalized marketing
  - Targeted ads based on browsing history
  - Recommend media based on previous consumption
- Healthcare
  - Diagnose illnesses based on symptoms
  - Identify risk factors/illness in healthy patients
- Fraud detection
  - Identify transactions that are out of character
- Predict outcome of sporting events
  - Google NCAA March Madness predictions on Kaggle
- Property valuation
  - Attributes of properties can be used to build models that predict possible future sales value


Before You Start

Step 1: Educate the Organization

- Understand the technology.
- What questions need answering?
- Evaluate business needs before developing solutions.

Step 2: Map Problems to Technology

- Lots of data ≠ need for ML
  - Data analysis
  - BI Tools/Excel
- Consider rule-based systems
- Is the problem one of:
  - Predicting outcomes?
  - Classifying objects?
  - Grouping objects?
- Use deep learning only when necessary

Step 2.5: Consider Humans

- Building ML based solutions are expensive
- The model may work, but the product may not.
- Substitute model for humans while tweaking output.

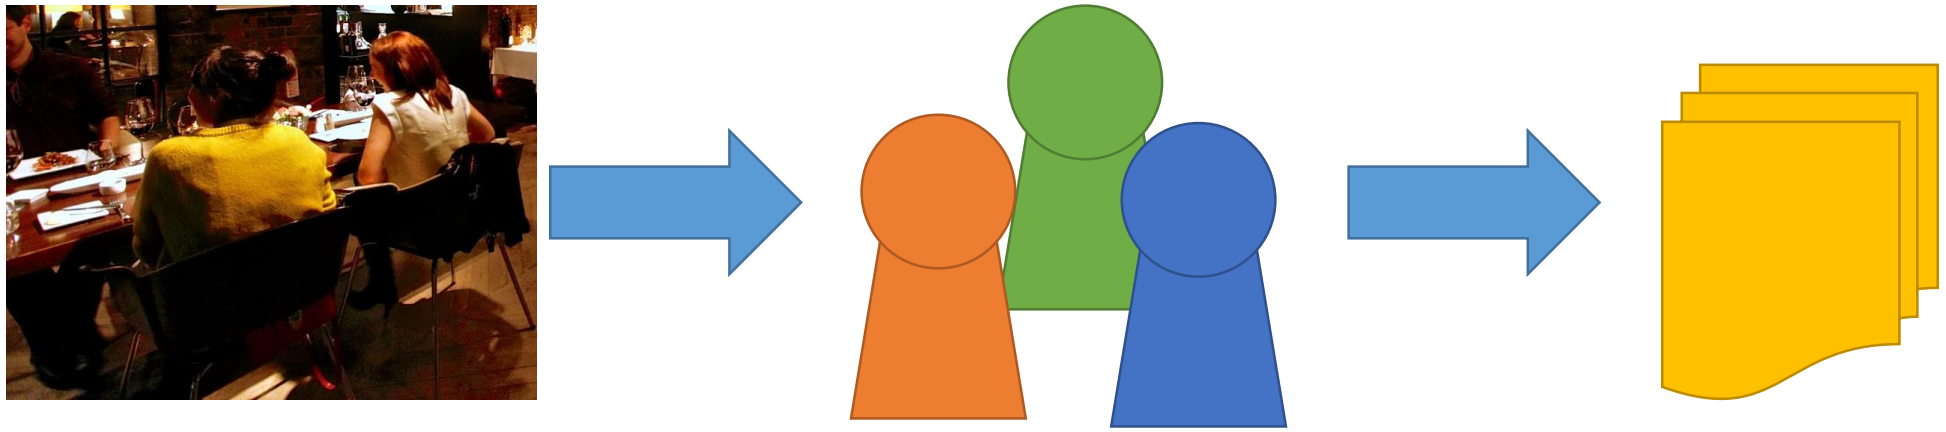

Building an AI/ML Solution

Building an AI/ML Solution

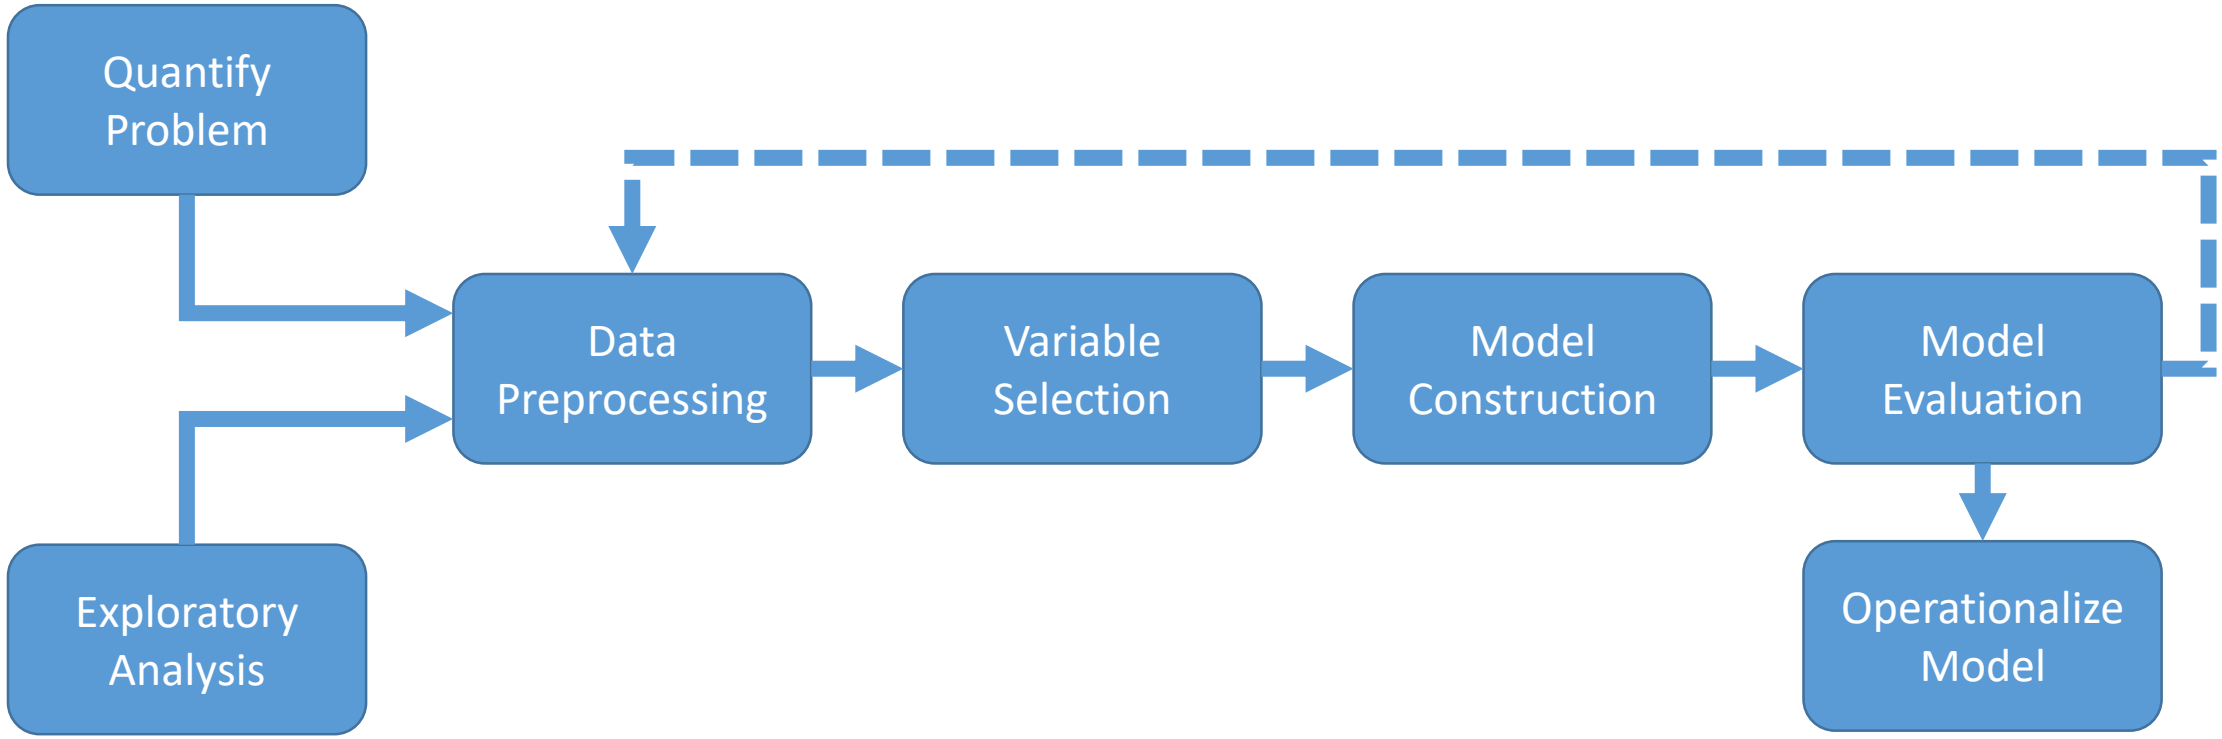

Quantify the Problem

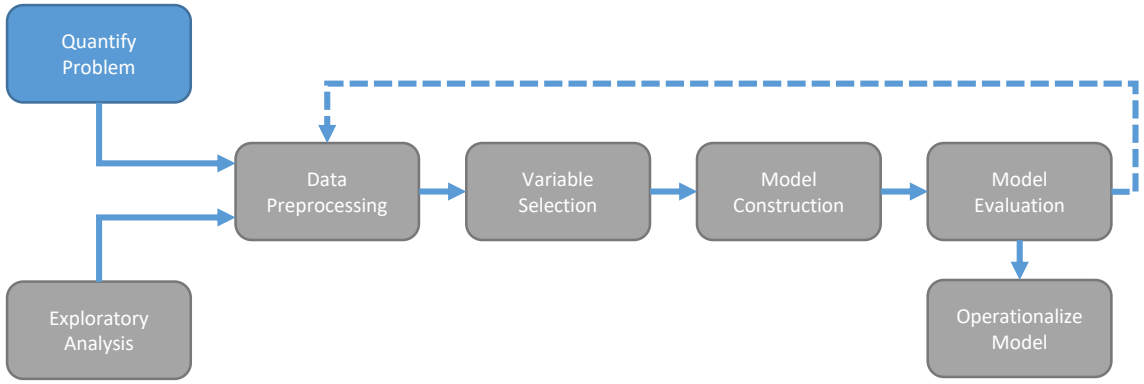

- Be very specific about what you want
- Words like “best” do not mean the same thing to everyone
- Business users and machine learning experts must communicate

Exploratory Analysis on Available Data

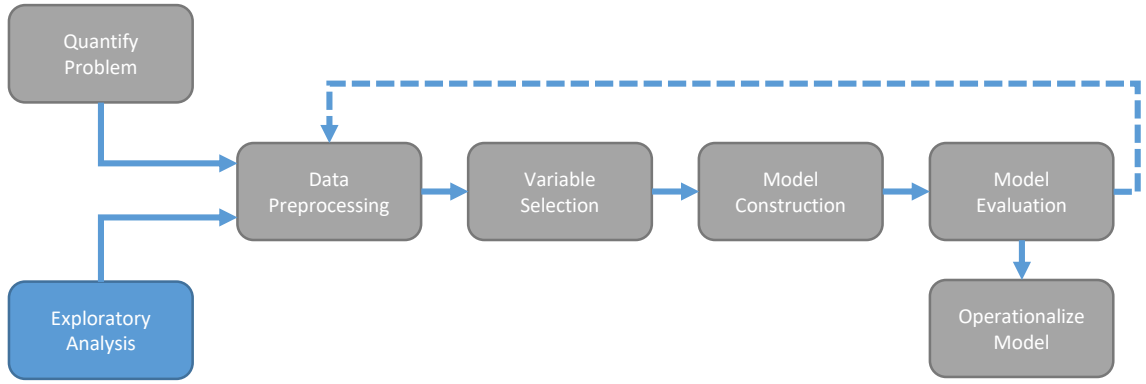

- What databases are available?
- How well are they documented?
- What is their purpose?
- What business processes are they involved in?
- What type of data do they contain?
- How are they currently used in decision making processes?

Preprocess Data

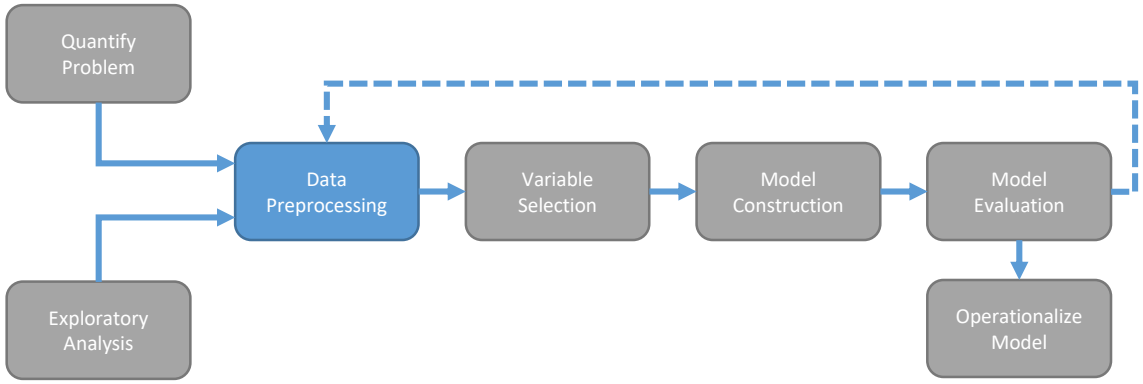

- The process of getting data ready for ML
- Creating new data from other data
- Data cleaning
- Consolidating data from different sources
- Converting between formats

Variable (Feature) Selection

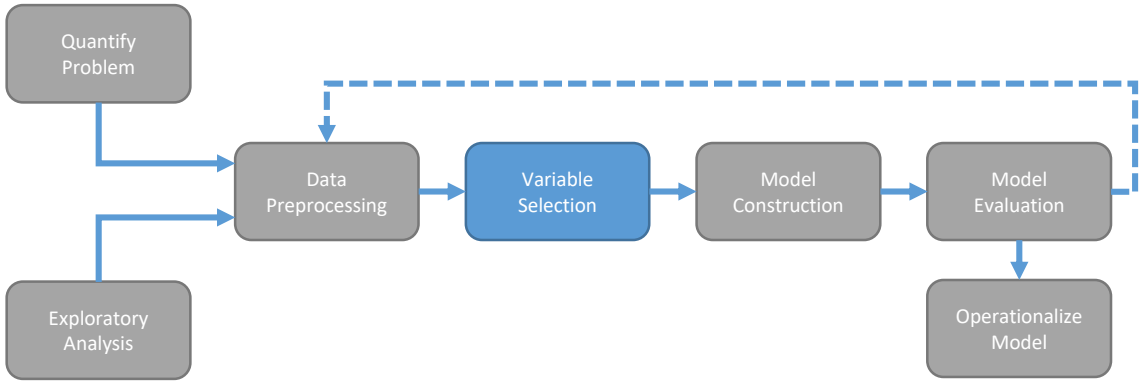

- Select relevant variables
- Discard irrelevant variables
- Leads to better results
- Shorter training times

Model Construction

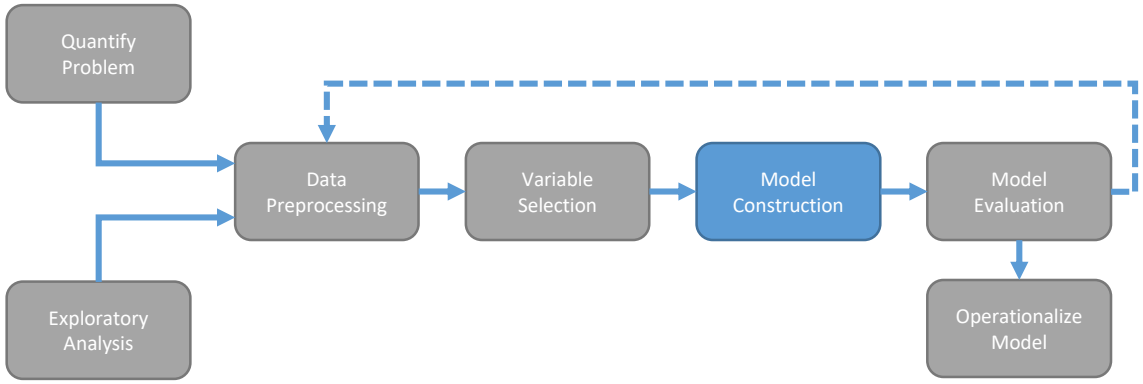

- Feed prepared data to candidate ML algorithms
- Set relevant (hyper)parameters
- Model is produced that can make predictions based on new inputs

Model Evaluation

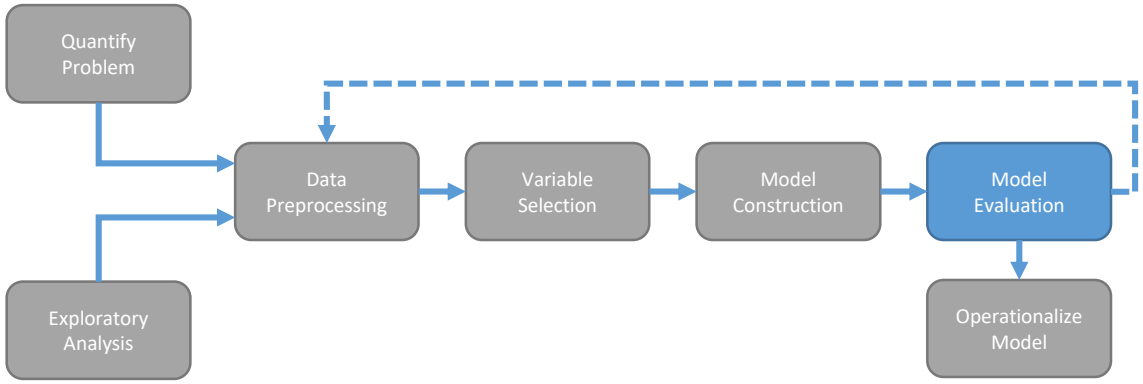

- Evaluation on accuracy
- More sophisticated evaluation methods also required
  - Sensitivity
  - Specificity
  - AUC - ROC Curve
  - Etc.
- Machine learning is an iterative process

Operationalize Model

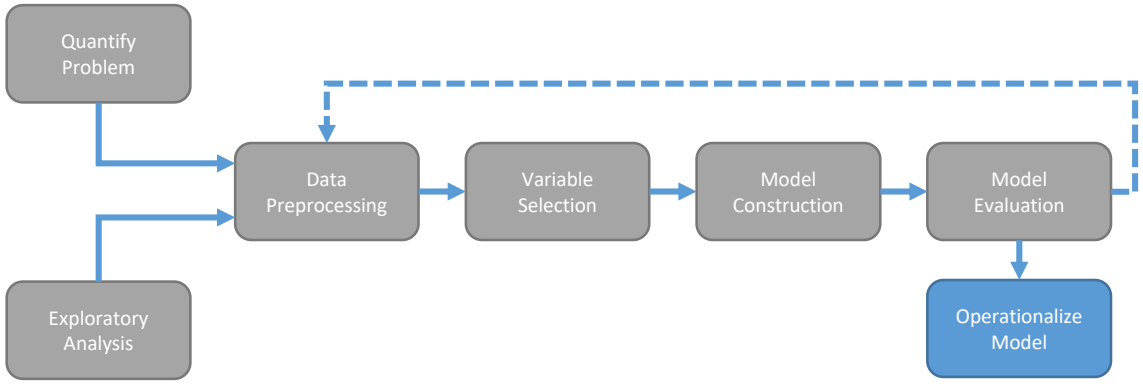

- Put the model to use
- Don’t forget to develop software around the model
- REST API
  - Single record
  - Batch of records
  - CSV
- Model output as input of another model

AI at Demand Solutions

- Review of our current forecasting methodology
- New forecasting engine
  - New methods
  - Machine Learning based forecasting
- Anomaly detection
  - Clean historical data
  - Provides for more accurate forecasts

AI at Demand Solutions

- Incorporation of external data sources – ML models
- Classification of imported data – automate importing
- Natural language based commands via Cortana
- Predicting machine failure using real-time data

- Microsoft Azure is a key component

Conclusion

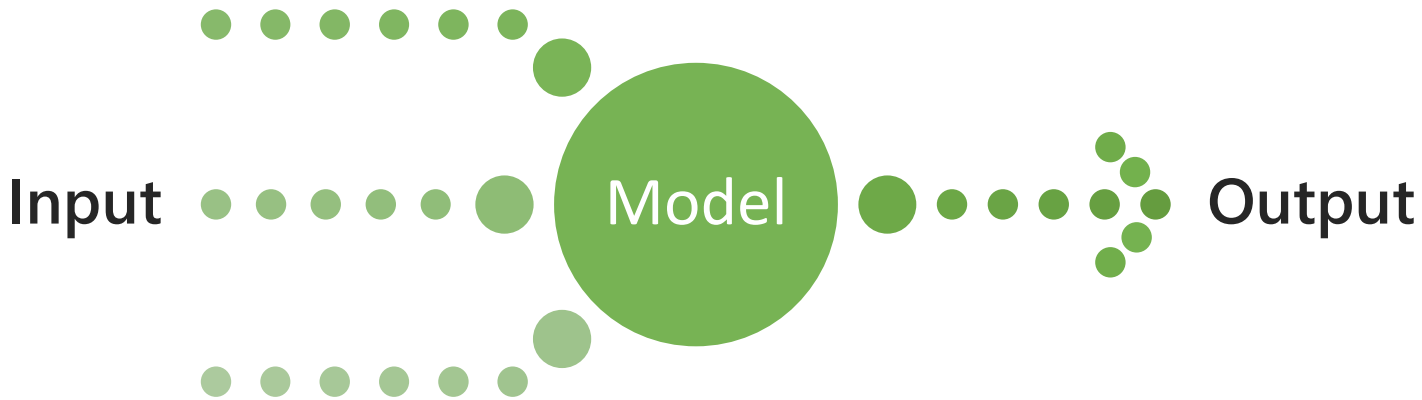
Input
Model
Output

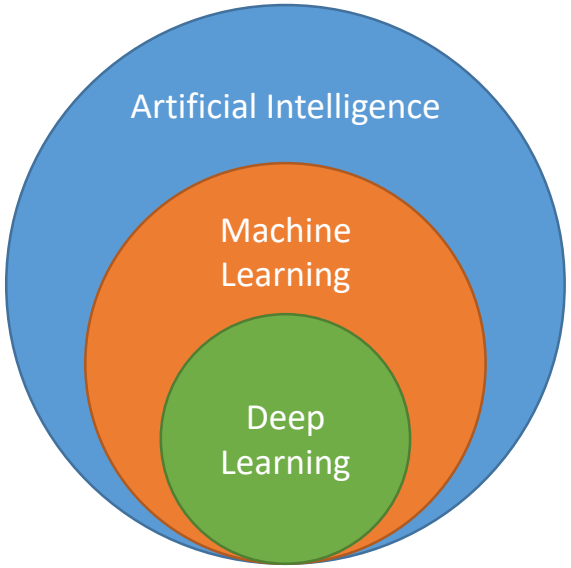
Artificial Intelligence
Machine Learning
Deep Learning

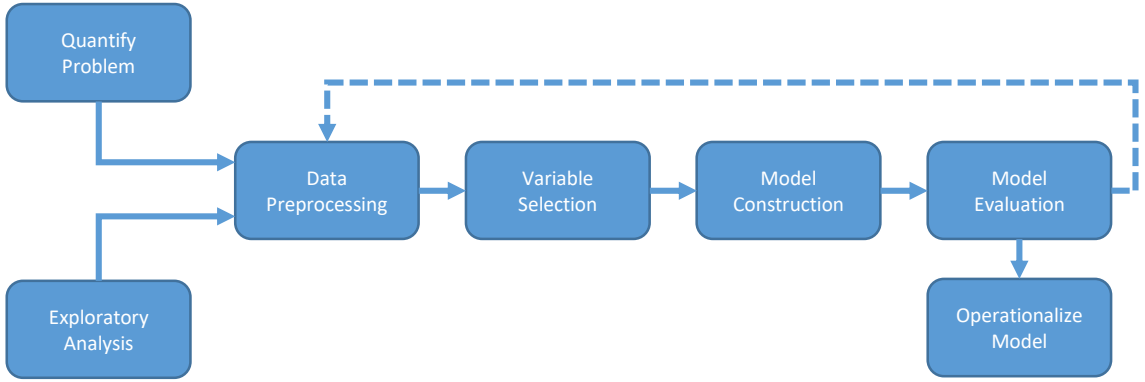
Quantify Problem
Exploratory Analysis
Data Preprocessing
Variable Selection
Model Construction
Model Evaluation
Operationalize Model

References

- Artificial Intelligence for Business: What You Need to Know about Machine Learning and Neural Networks, Doug Rose
  https://www.amazon.com/Artificial-Intelligence-Business-Learning-Networks-ebook/dp/B07CCLXFY1
- Artificial Intelligence and Machine Learning for Business: A No-Nonsense Guide to Data Driven Technologies, Steven Finlay
  https://www.amazon.com/Artificial-Intelligence-Machine-Learning-Business-ebook/dp/B0719VH2KR
- Data Jujitsu: The art of turning data into product, DJ Patil
  https://www.amazon.com/Data-Jujitsu-Turning-into-Product-ebook/dp/B008HMN5BE
- Hands-On Machine Learning with Scikit-Learn and TensorFlow, Aurélien Géron
  https://www.amazon.com/Hands-Machine-Learning-Scikit-Learn-TensorFlow/dp/1491962291
- The State of Machine Learning Adoption in the Enterprise, Ben Lorica & Paco Nathan
  https://www.oreilly.com/library/view/the-state-of/9781492044383/
- History of Artificial Intelligence - Wikipedia
  https://en.wikipedia.org/wiki/History_of_artificial_intelligence


QUESTIONS?


THANK YOU