Alberta
Machine
Intelligence
Institute

Alliance Pipeline Seminar

The power and pitfalls of machine learning in supply chain management

Anna Koop
Senior Scientific Advisor, Amii

June 5, 2019

How would you describe machine learning?

Who is Amii?

Vision

Inspire world-changing intelligence for good and for all.

Make Alberta the premier destination for machine intelligence research, development, and commercialization.

Amii Researchers

Michael Bowling

Alona Fyshe

Randy Goebel

Russ Greiner

Robert Holte

Patrick M. Pilarski

Dale Schuurmans

Or Sheffet

Richard S. Sutton

Csaba Szepesvári

Martha White

James Wright

Yutaka Yasui

Osmar Zaïane

Angel Chang

Kevin Leyton-Brown

Mark Schmidt

Four focus areas

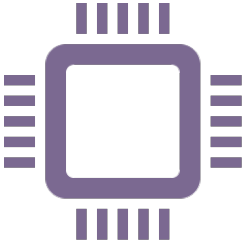

Supporting world-class research and training

Growing machine intelligence capacity in business

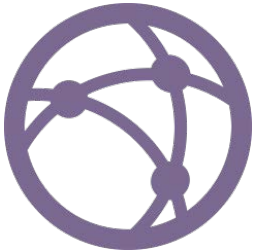

Connecting global innovators

Boosting machine intelligence literacy in business

What to expect from this seminar

What is ML?
Understanding definitions

What can it do?
Examples related to supply chain

How to succeed
Projects need expertise

Case study
Grocery store competition

What is ML?

Understanding definitions

What is Artificial Intelligence?

AI Example: Pathfinding

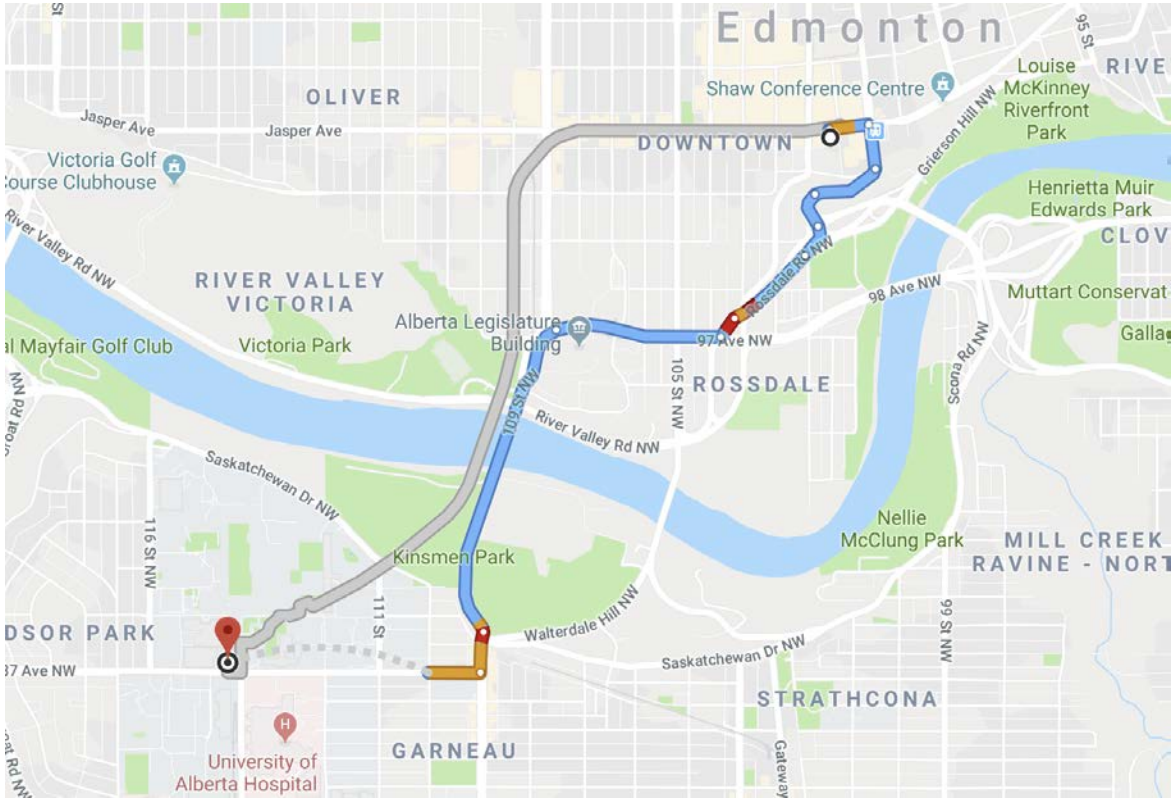

AI Example: ‘Smart’ Assistant

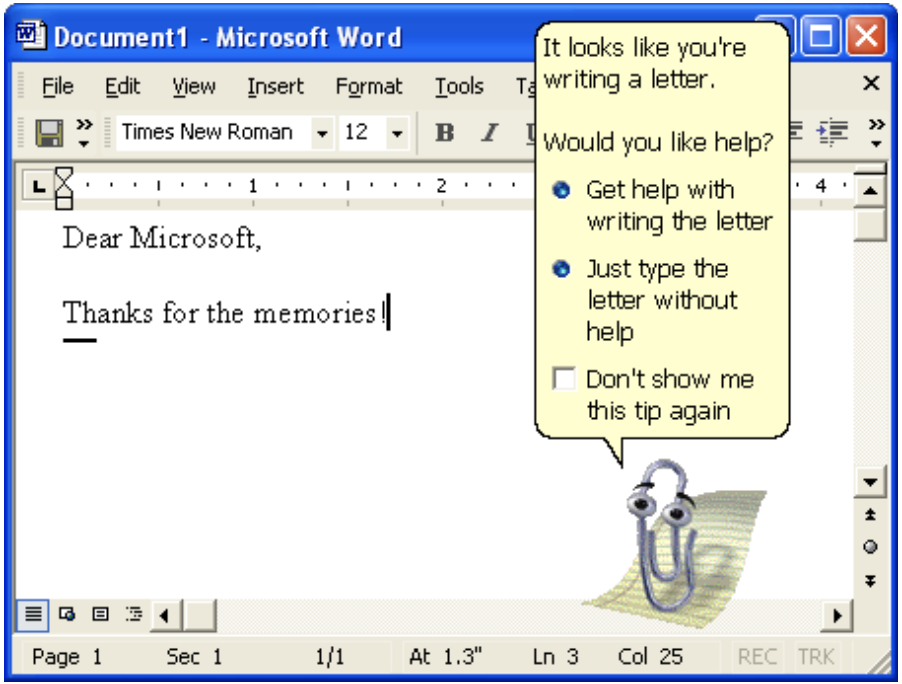

https://www.artsy.net/article/artsy-editorial-life-death-microsoft-clippy-paper-clip-loved-hate

AI Example: “Autonomous” Robots

https://en.wikipedia.org/wiki/Furby

By ITU Pictures CC BY 2.0,
https://commons.wikimedia.org/w/index.php?curid=63678155

AI Example: Autonomous Cars

http://www.darpa.mil/grandchallenge05/high_res

https://www.theguardian.com/technology/2014/may/28/google-self-driving-car-how-does-it-work

What is Machine Learning?

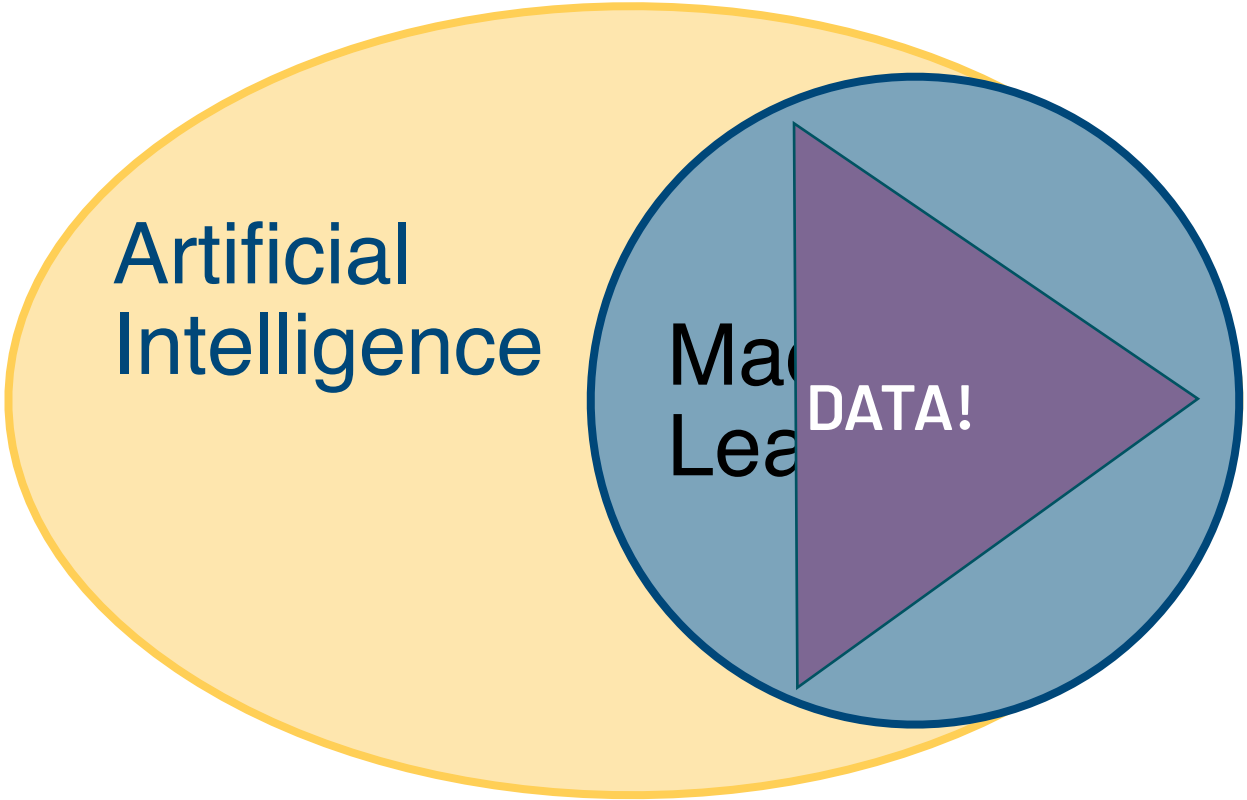

Types of Machine Learning

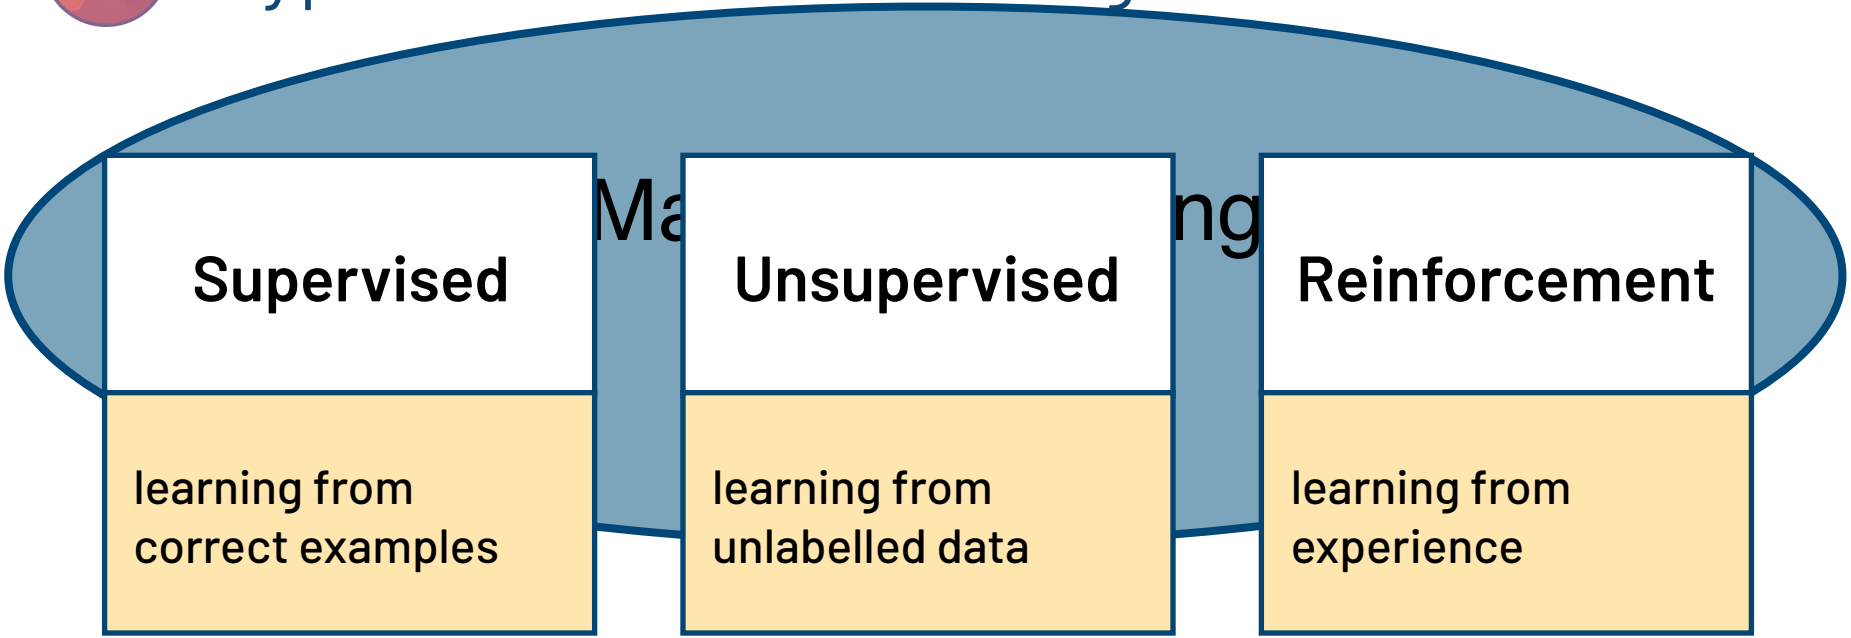

What is AI - ML?

Classic Artificial Intelligence

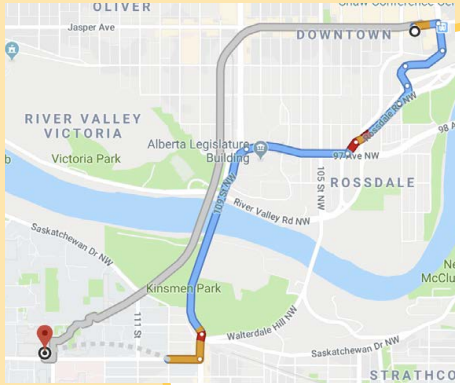

The Output of ML

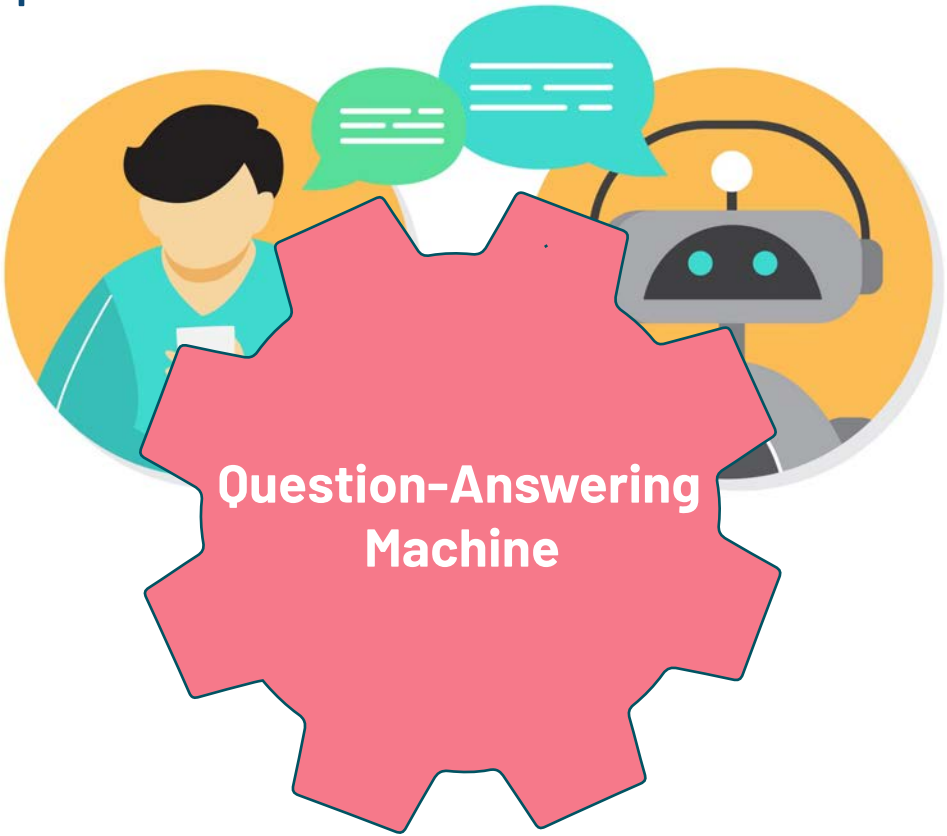

Machine Learning Process

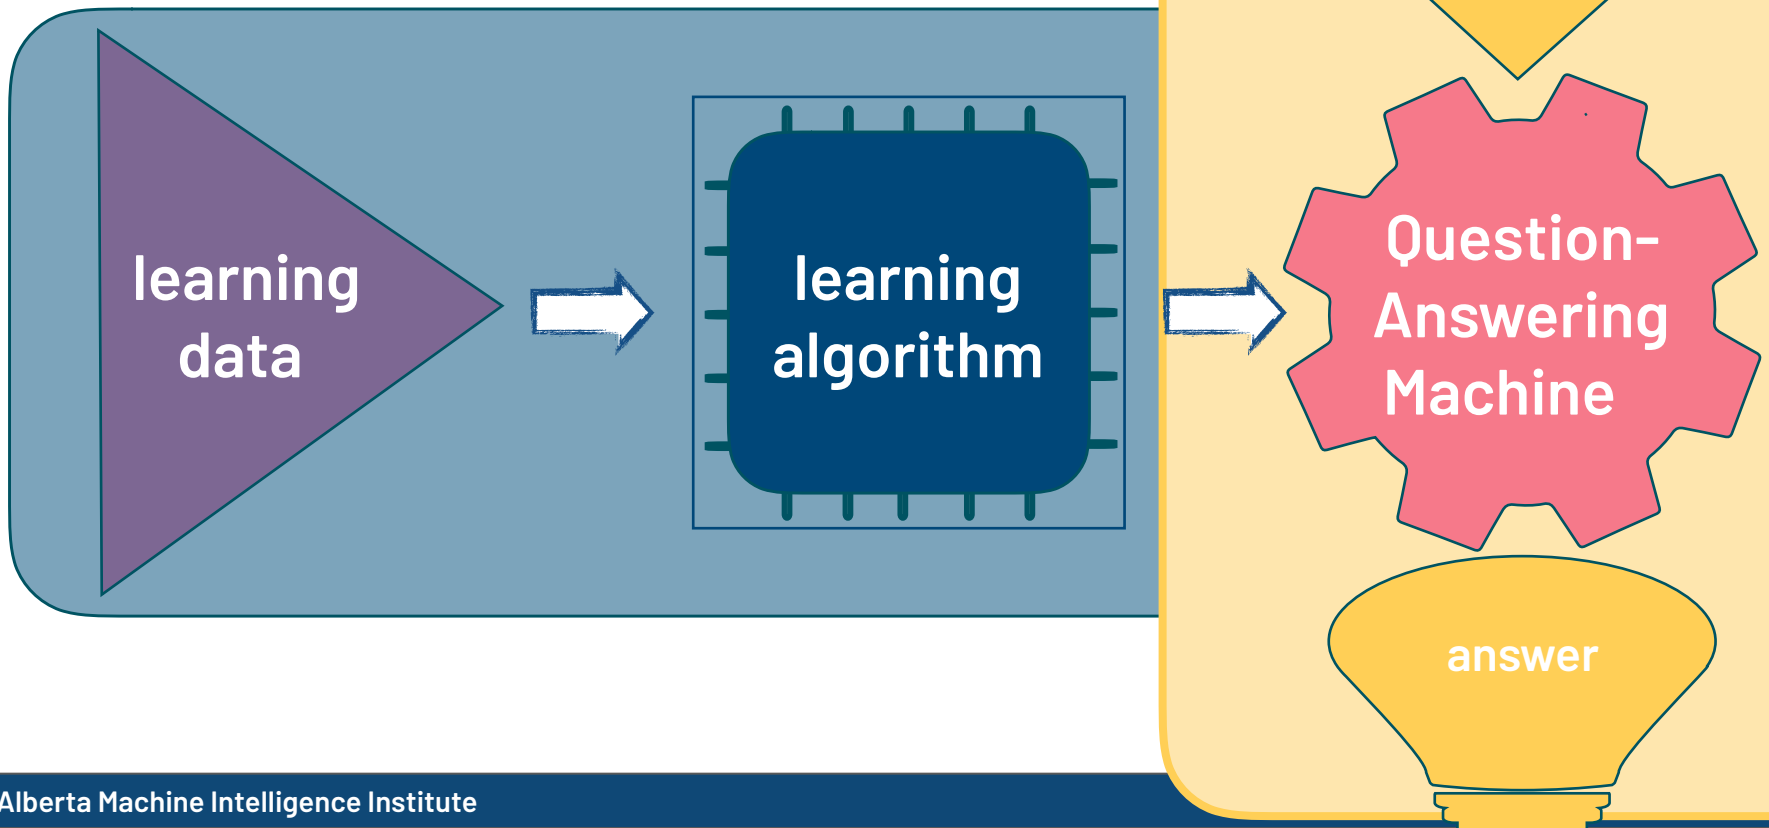

Data Science

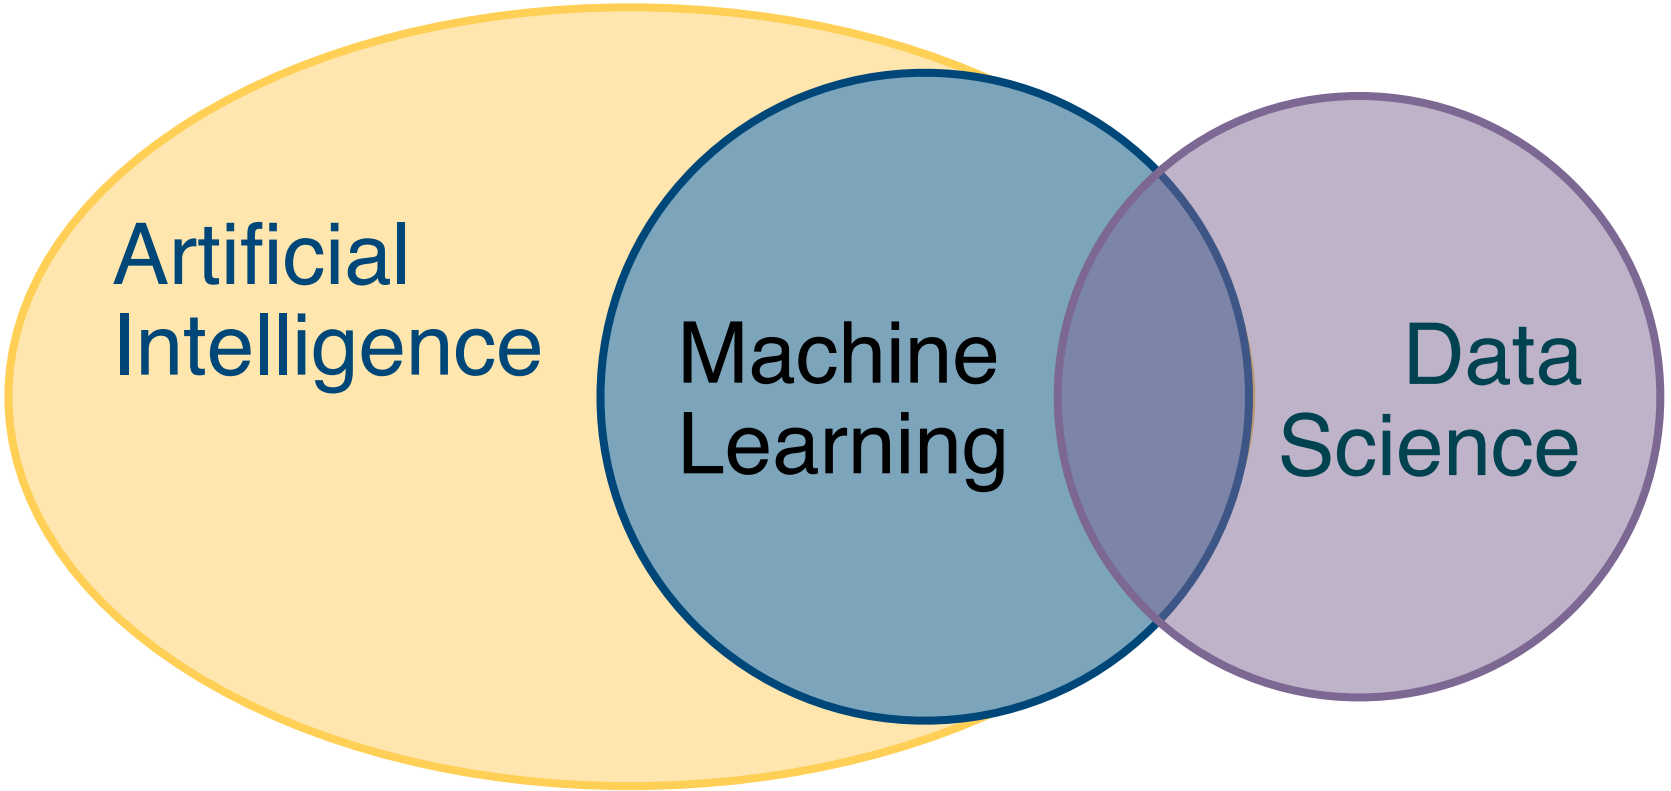

What is Deep Learning?

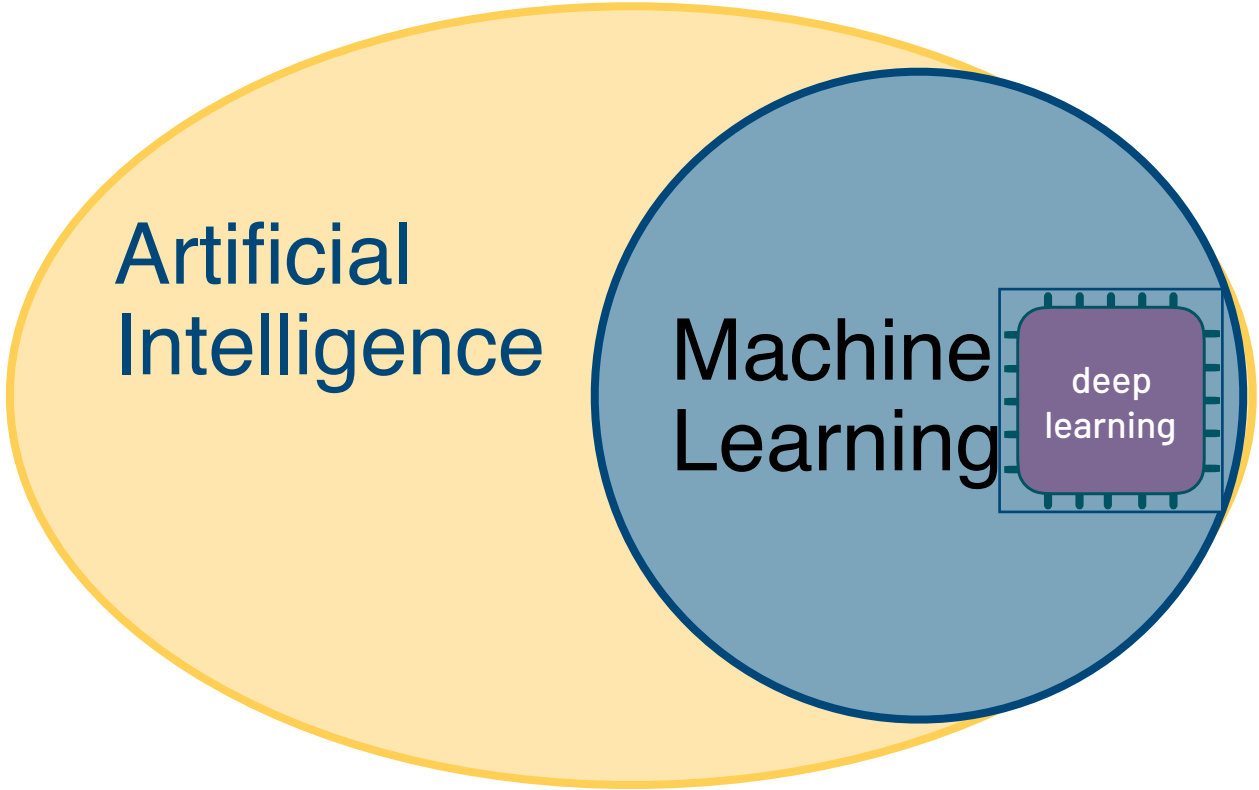

Local Deep Learning Successes

About those Questions...

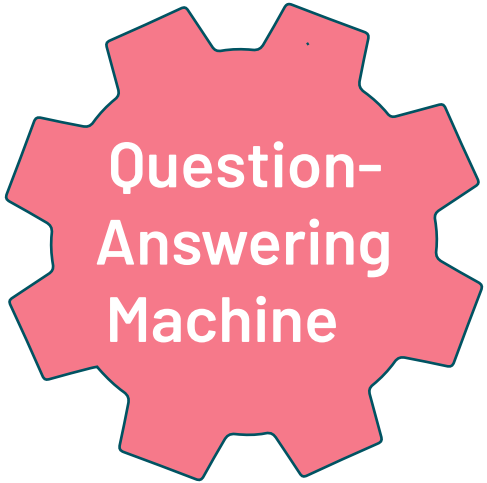

Subtasks

- Detect and classify objects
- Generate valid sentences
- Rank sentences

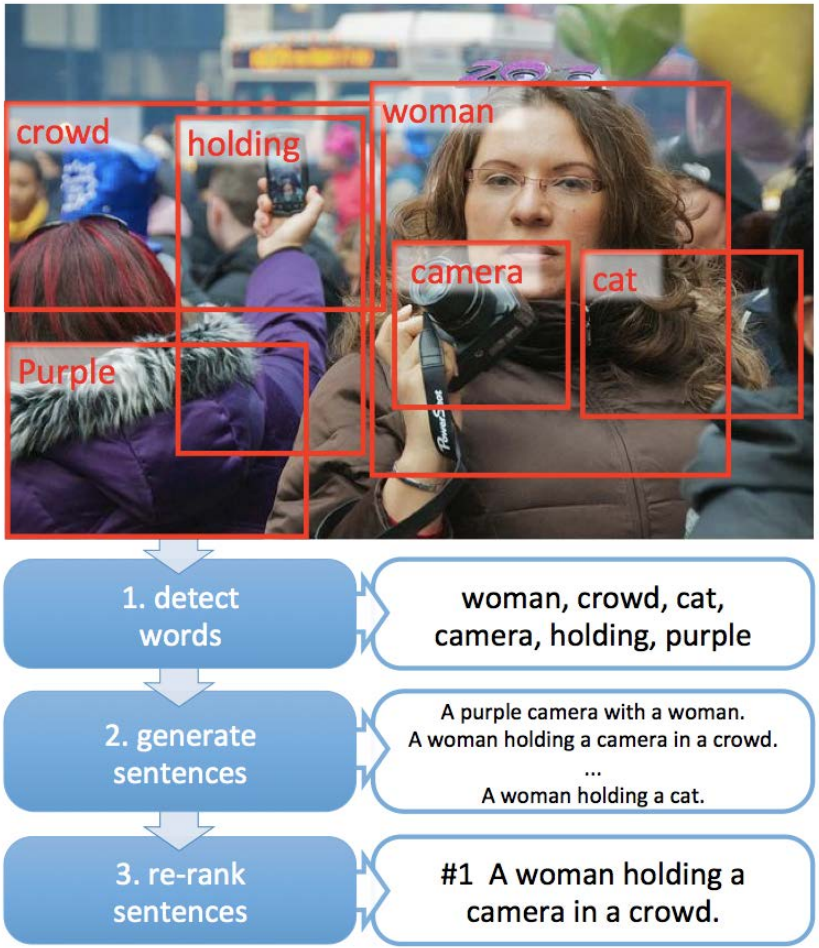

https://blogs.technet.microsoft.com/machinelearning/2014/11/18/rapid-progress-in-automatic-image-captioning/

Image captioning results

"man in black shirt is playing guitar."

"construction worker in orange safety vest is working on road."

"two young girls are playing with lego toy."

https://towardsdatascience.com/image-captioning-in-deep-learning-9cd23fb4d8d2

Image captioning...?

http://gizmodo.com/this-neural-networks-hilariously-bad-image-descriptions-1730844528

Fooling a Neural Network

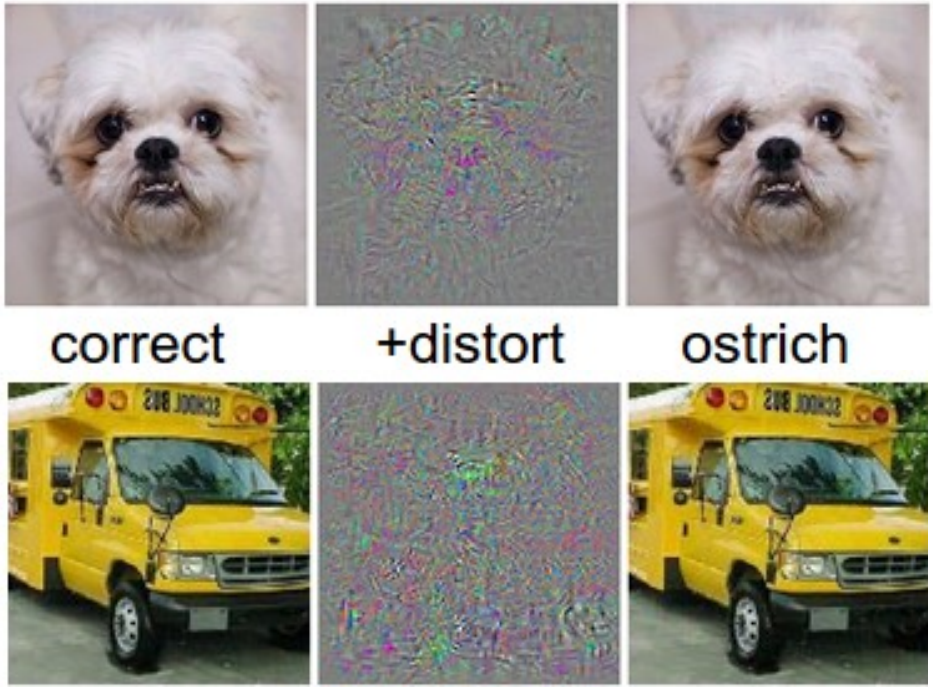

http://karpathy.github.io/2015/03/30/breaking-convnets/

https://arxiv.org/abs/1312.6199

Ceci n’est pas une banana

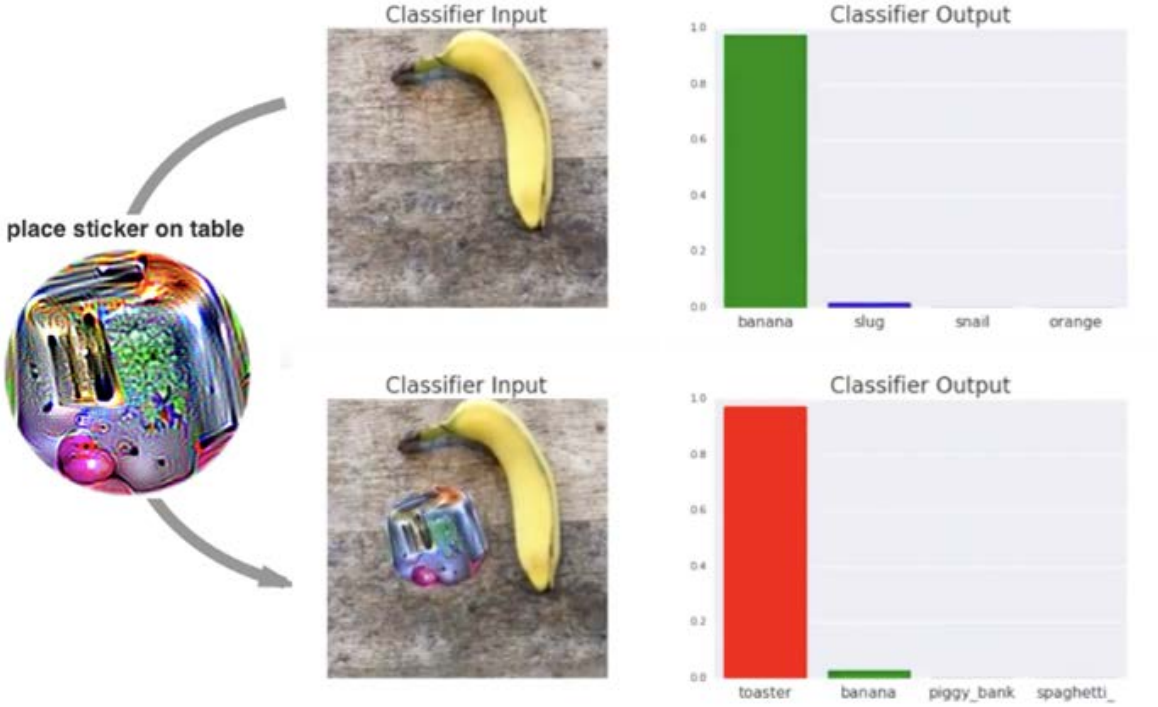

https://gizmodo.com/this-simple-sticker-can-trick-neural-networks-into-thin-1821735479
Image: Tom B. Brown/Dandelion Mané, https://arxiv.org/pdf/1712.09665.pdf

Essence of an Image

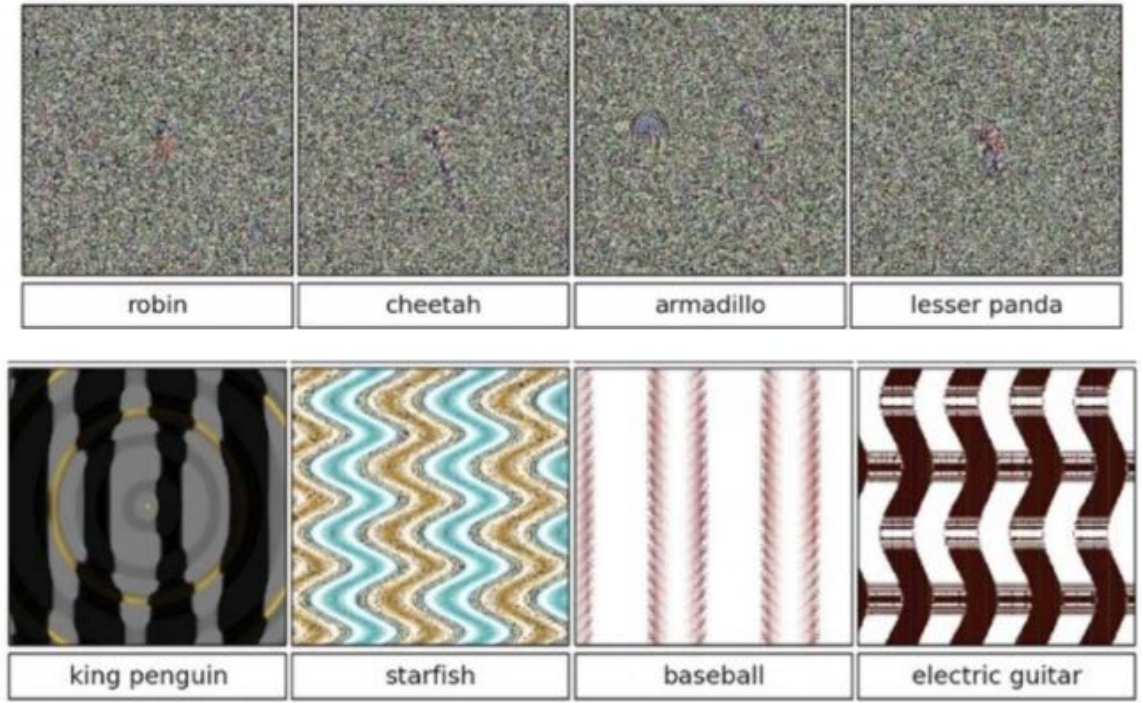

http://karpathy.github.io/2015/03/30/breaking-convnets/

https://arxiv.org/abs/1412.1897

Generalization is hard!

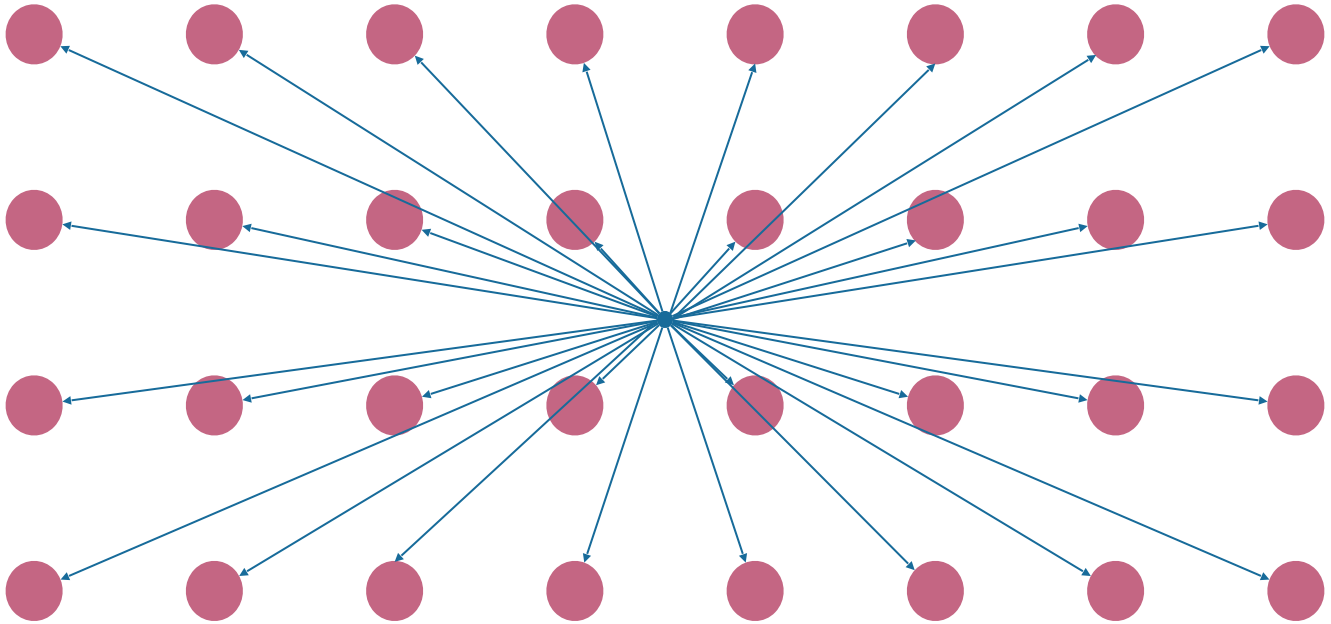

Focus on a specific part of the problem

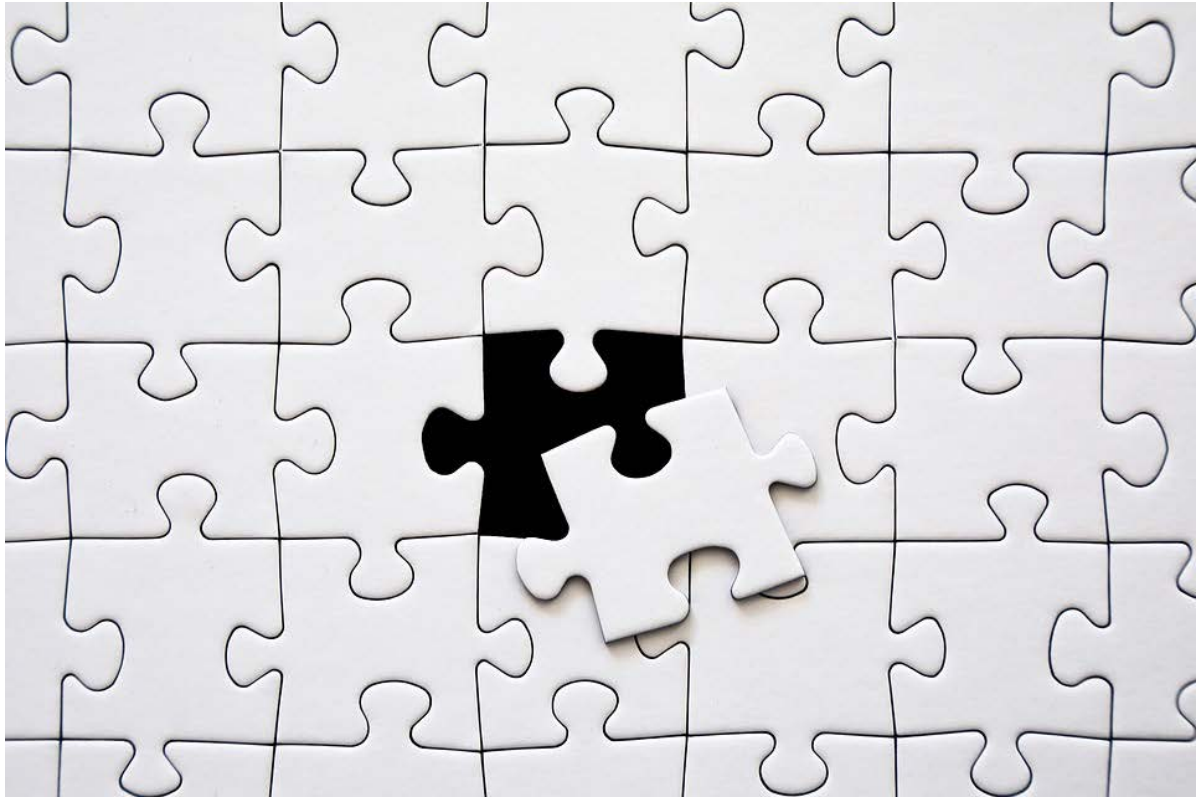

What can ML do?

Examples related to supply chain

Prediction (Machine Learning)

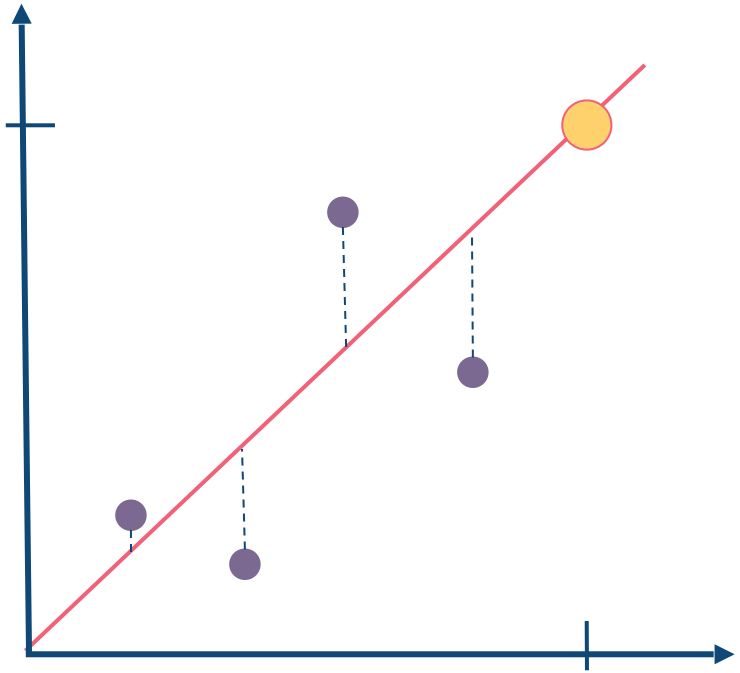

Inferring something that hasn't happened yet, or that you can't measure directly

Prediction (Machine Learning)

number of sales

Prediction (Machine Learning)

Daniel Case [CC BY-SA 3.0 (https://creativecommons.org/licenses/by-sa/3.0)]

time to machine failure

Prediction (Machine Learning)

simulation / digital twins

https://media.giphy.com/media/3oz8xTfD5SrkAwNNUQ/source.gif

Classification (Machine Learning)

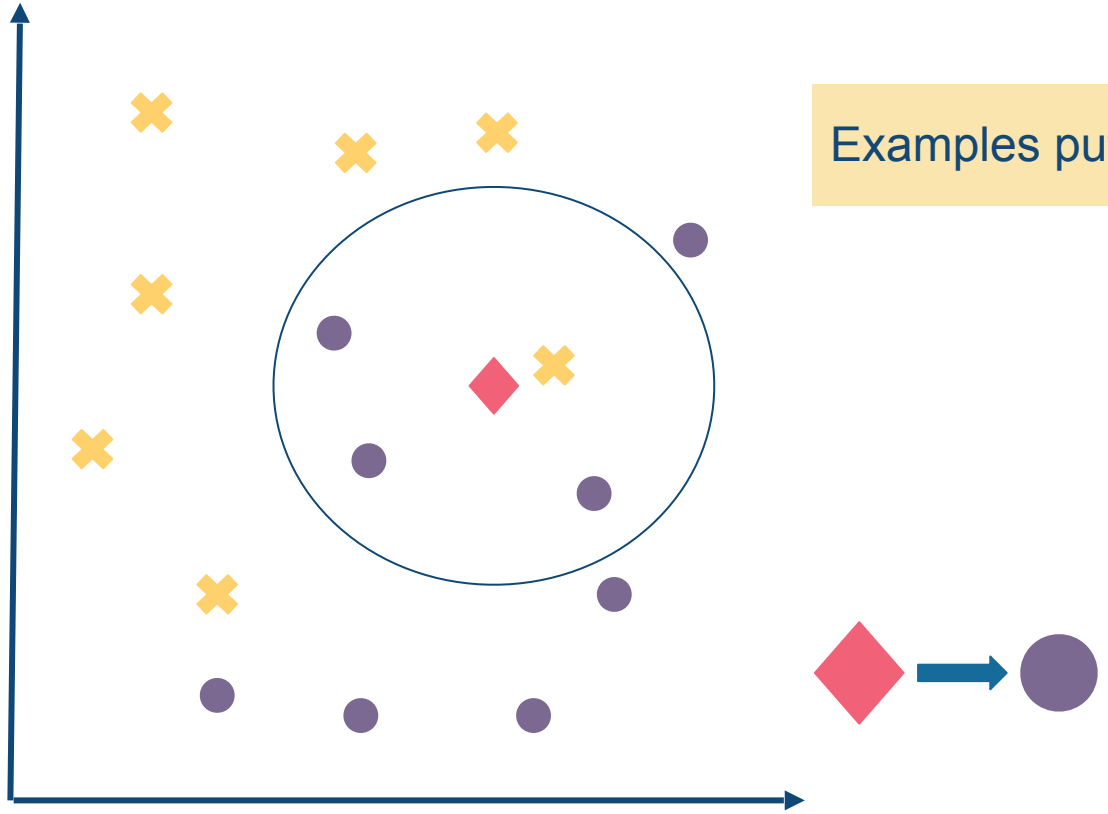

Examples put into categories

Classification (Machine Learning)

to buy or not to buy?

Classification (Machine Learning)

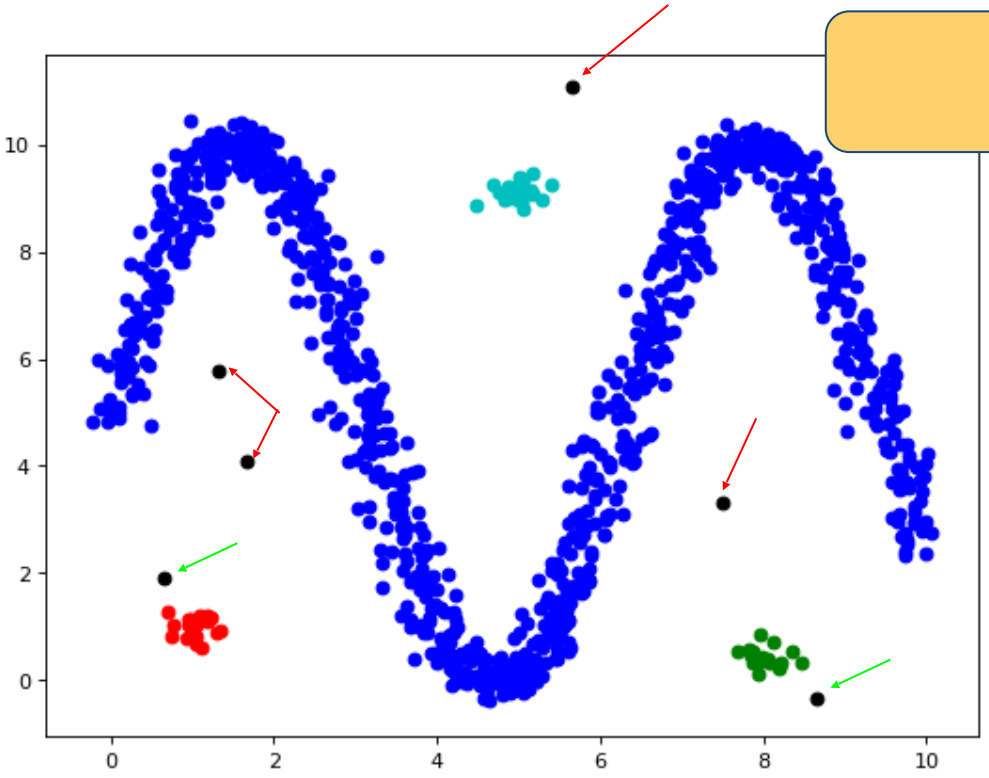

fraud detection

Local Outlier

Global Outlier

Classification (Machine Learning)

on-time delivery

Association (Machine Learning)

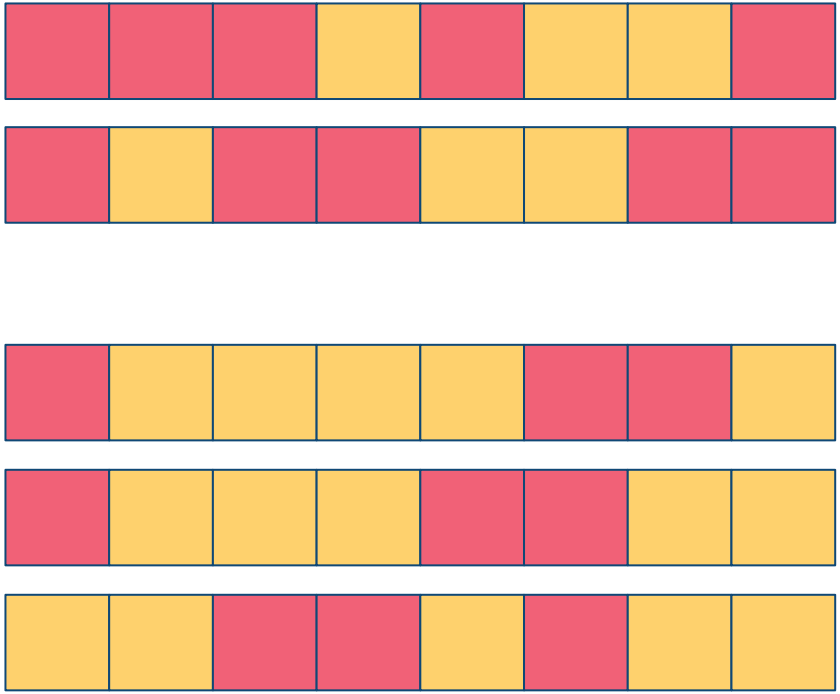

Finding links between people or things

Grouping people/items based on associations

Association (Machine Learning)

identify customer communities

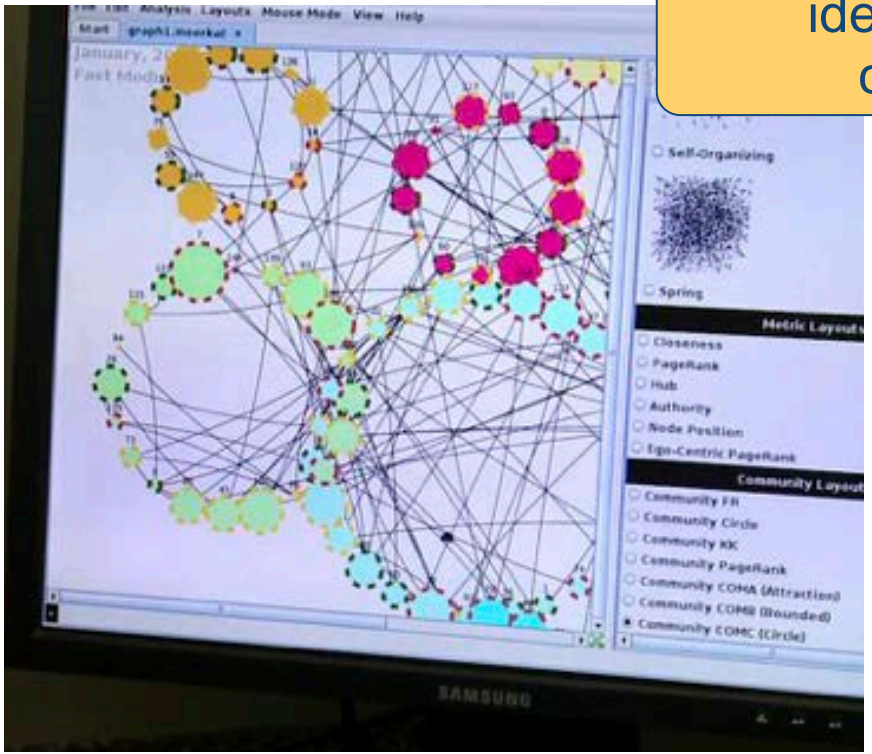

Association (Machine Learning)

market basket analytics

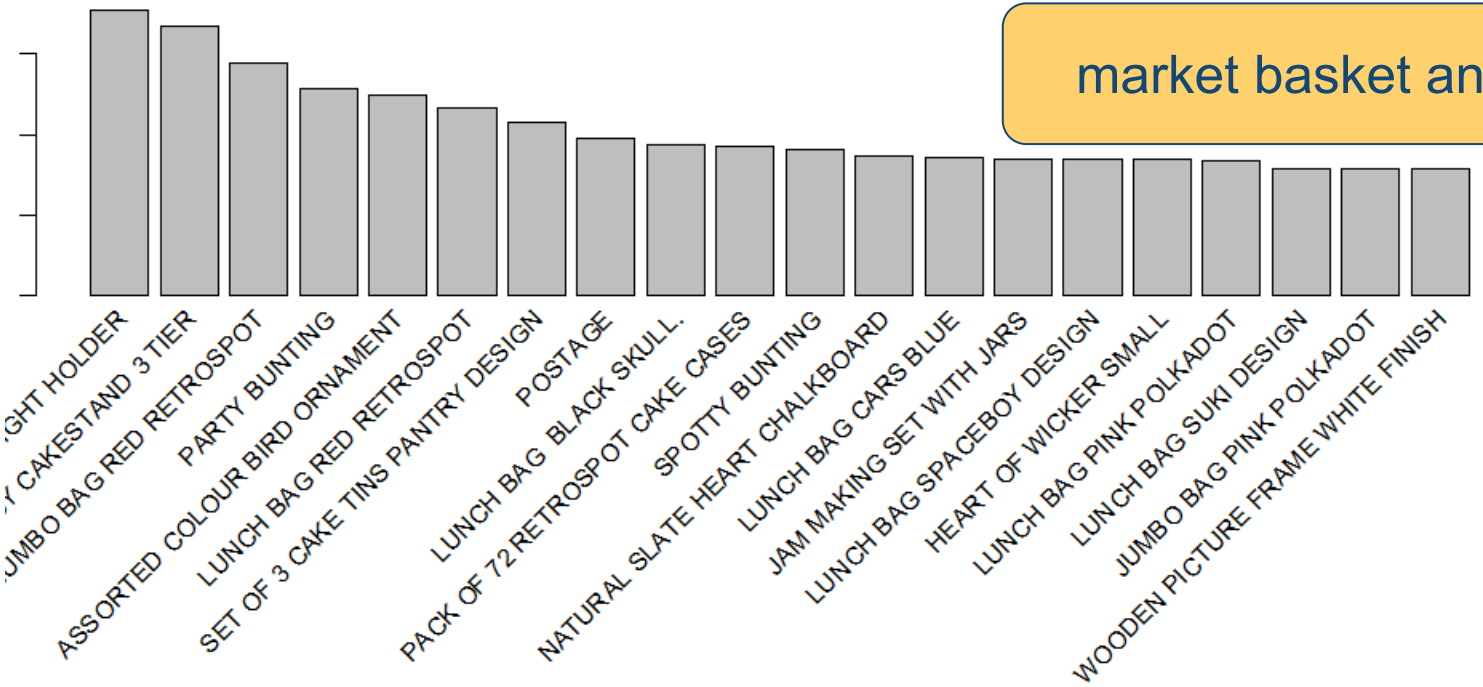

Figure 4. A bar plot of the support of the 20 most frequent items bought.

https://towardsdatascience.com/a-gentle-introduction-on-market-basket-analysis-association-rules-fa4b986a40ce

Association (Machine Learning)

recommender systems

amazon

Spotify®

Association (Machine Learning)

finding good suppliers

Creativity103 [CC BY 2.0 (https://creativecommons.org/licenses/by/2.0)]

Action Control (Machine Learning)

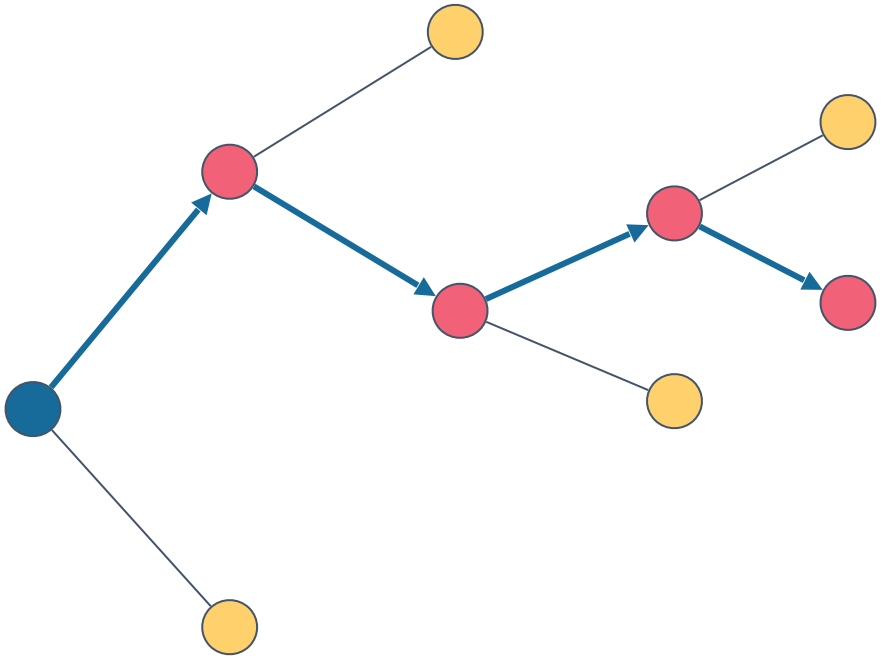

Choosing an action

Choice is made in response to the environment

Action Control (Machine Learning)

when to make an offer?

Action Control (Machine Learning)

when to contact clients?

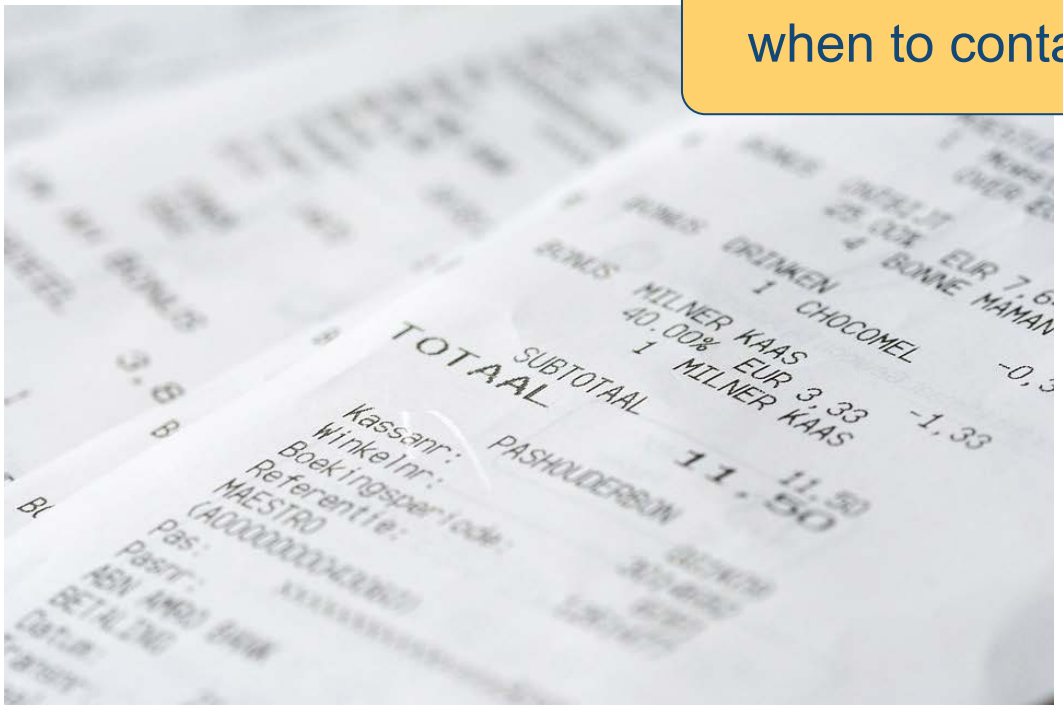

Action Control (Machine Learning)

robo warehouse

Planning (Classic AI)

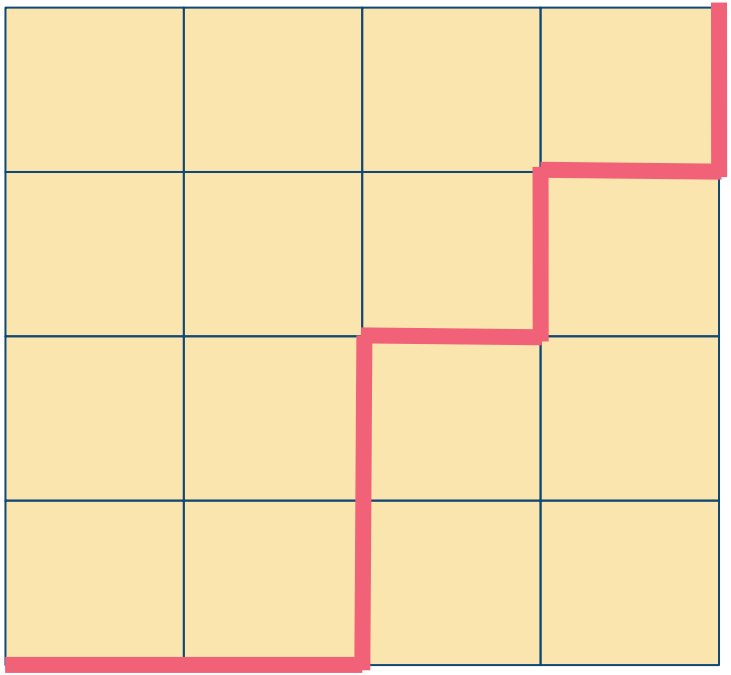

Finding the optimal path in a system with perfect information

Planning (Classic AI)

multi-stop route finding

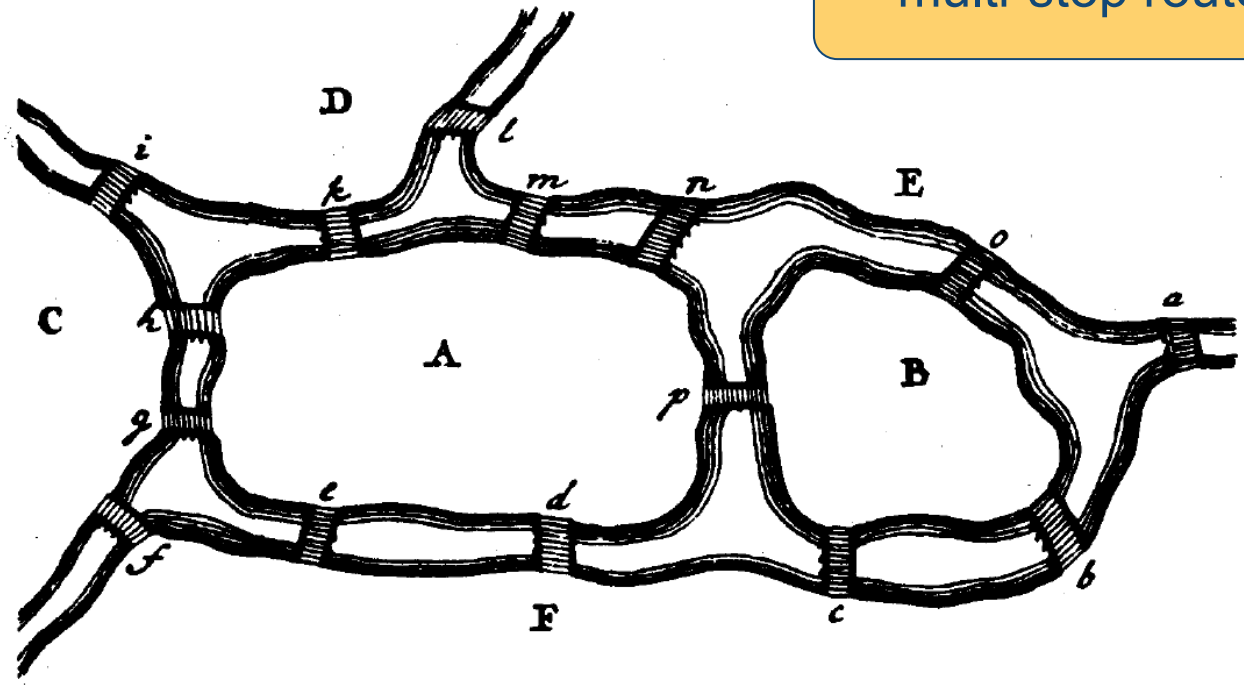

Leonhard Euler [Public domain]

Planning (Classic AI)

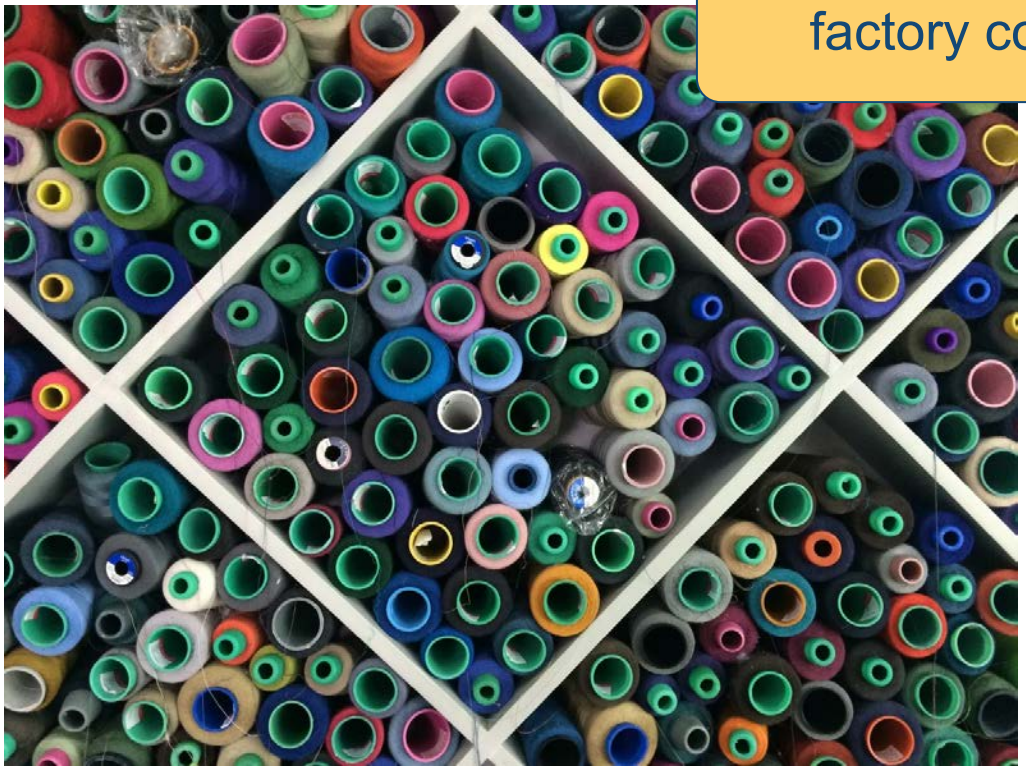

factory configuration

Planning (Classic AI)

warehouse configuration

How to succeed?

Projects need expertise

Data is the new oil

Incredibly valuable!

(but utterly useless in its default state)

Enables profound new technologies!

(but can create disasters if used poorly)

Data is not knowledge

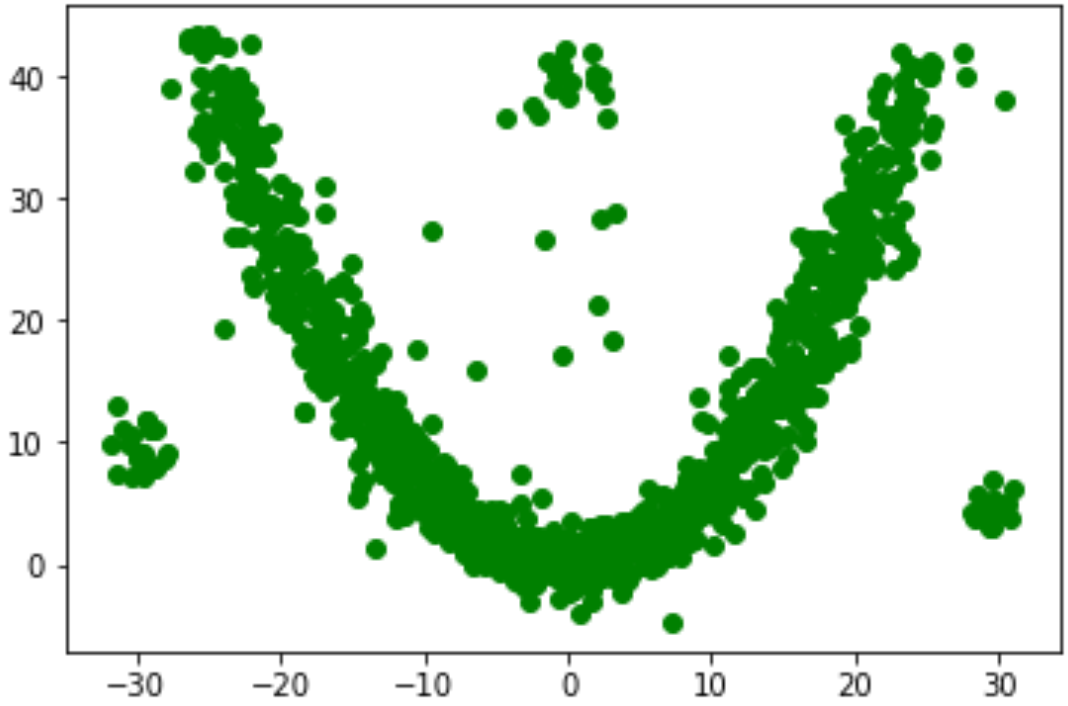

Data is not knowledge

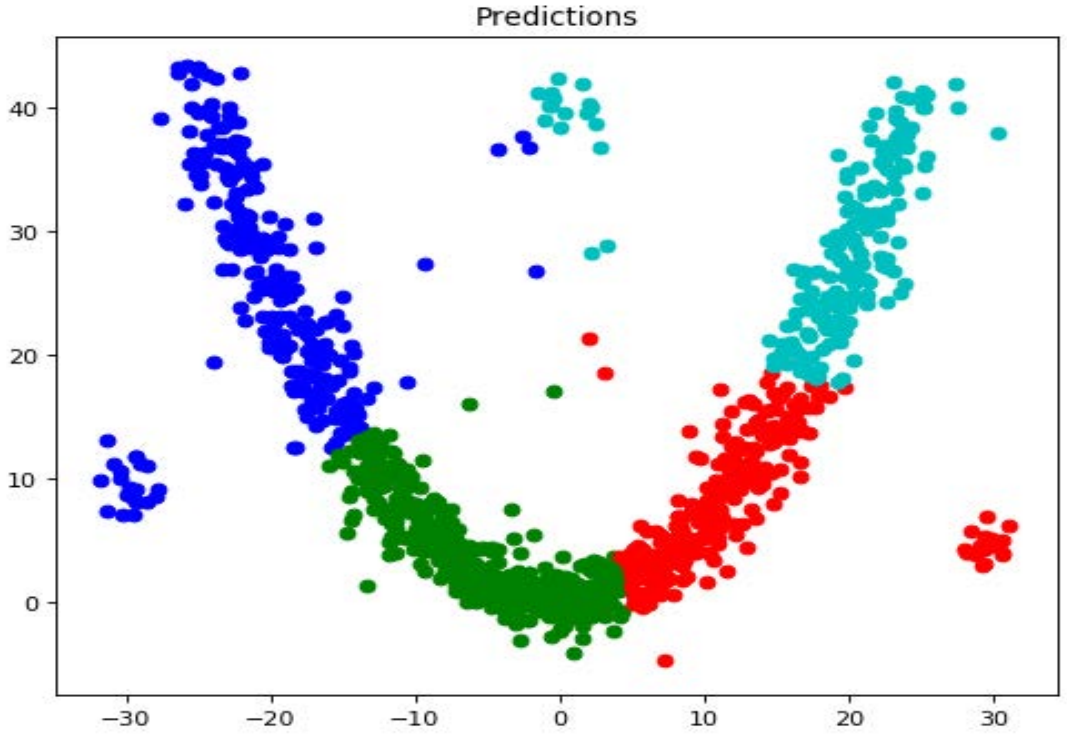

Garbage In, Garbage Out

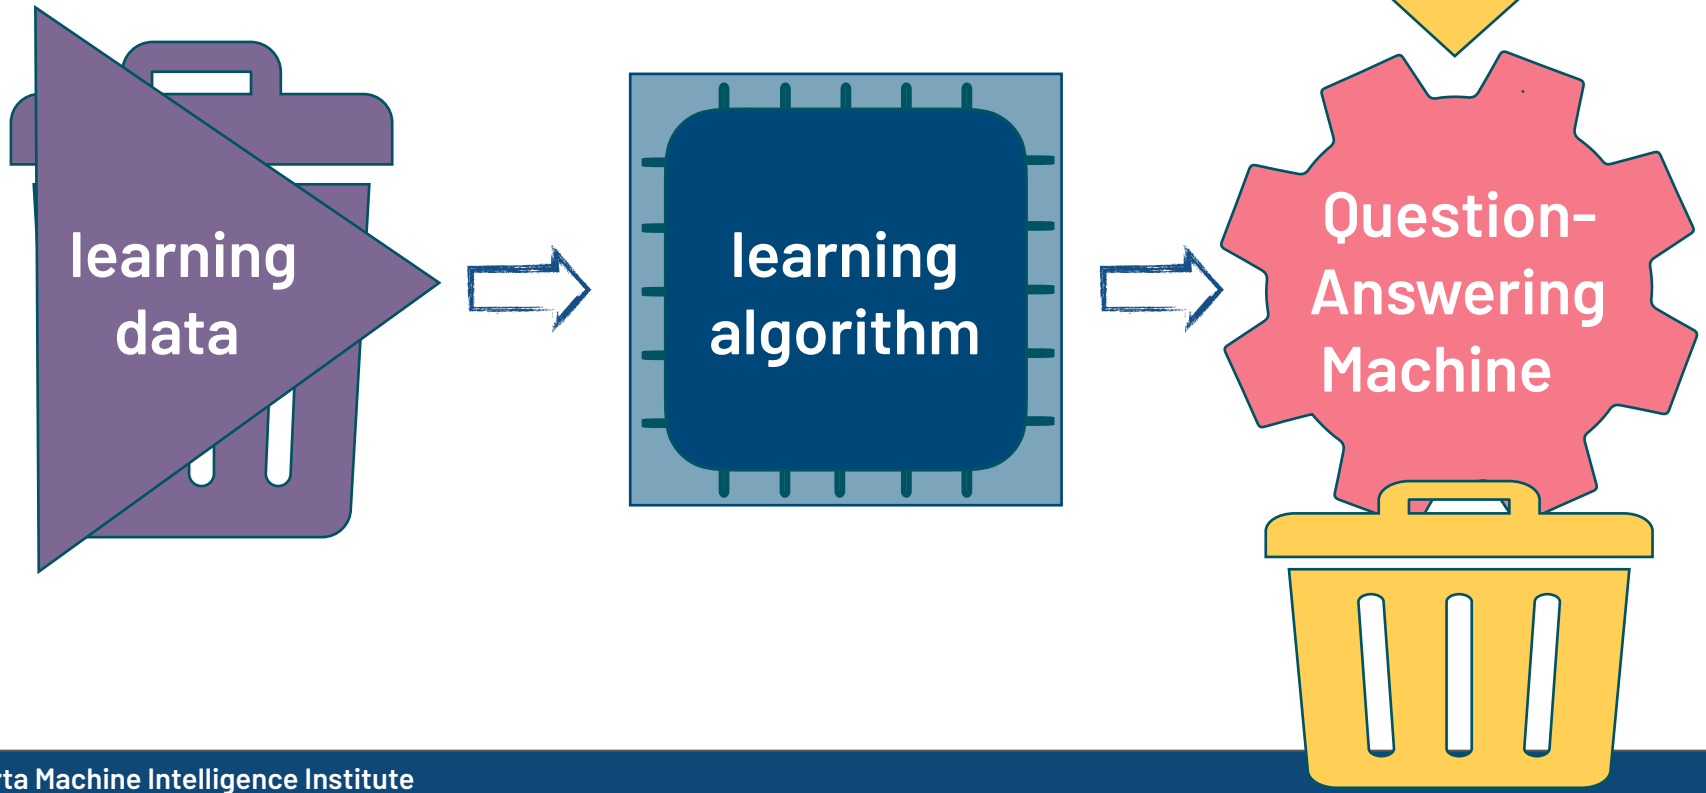

What is “garbage data”?

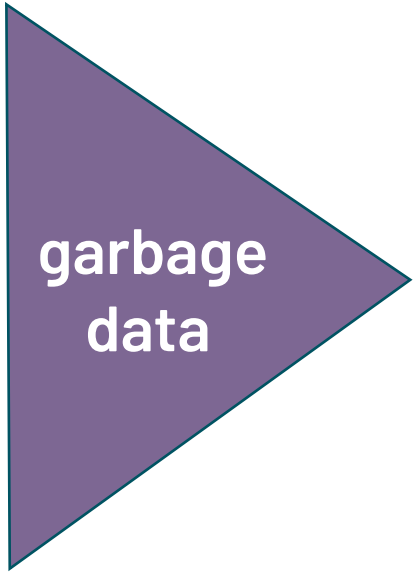

- Not enough examples
- Not enough variety
- Not enough detail
- Too many irrelevant details
- About the wrong aspect
- Different scales and sources

ML means answering questions...

Faster!

(eventually)

Cheaper!

(in operation)

More accurately!

(in the right settings)

More consistently!

Always!

When to automate?

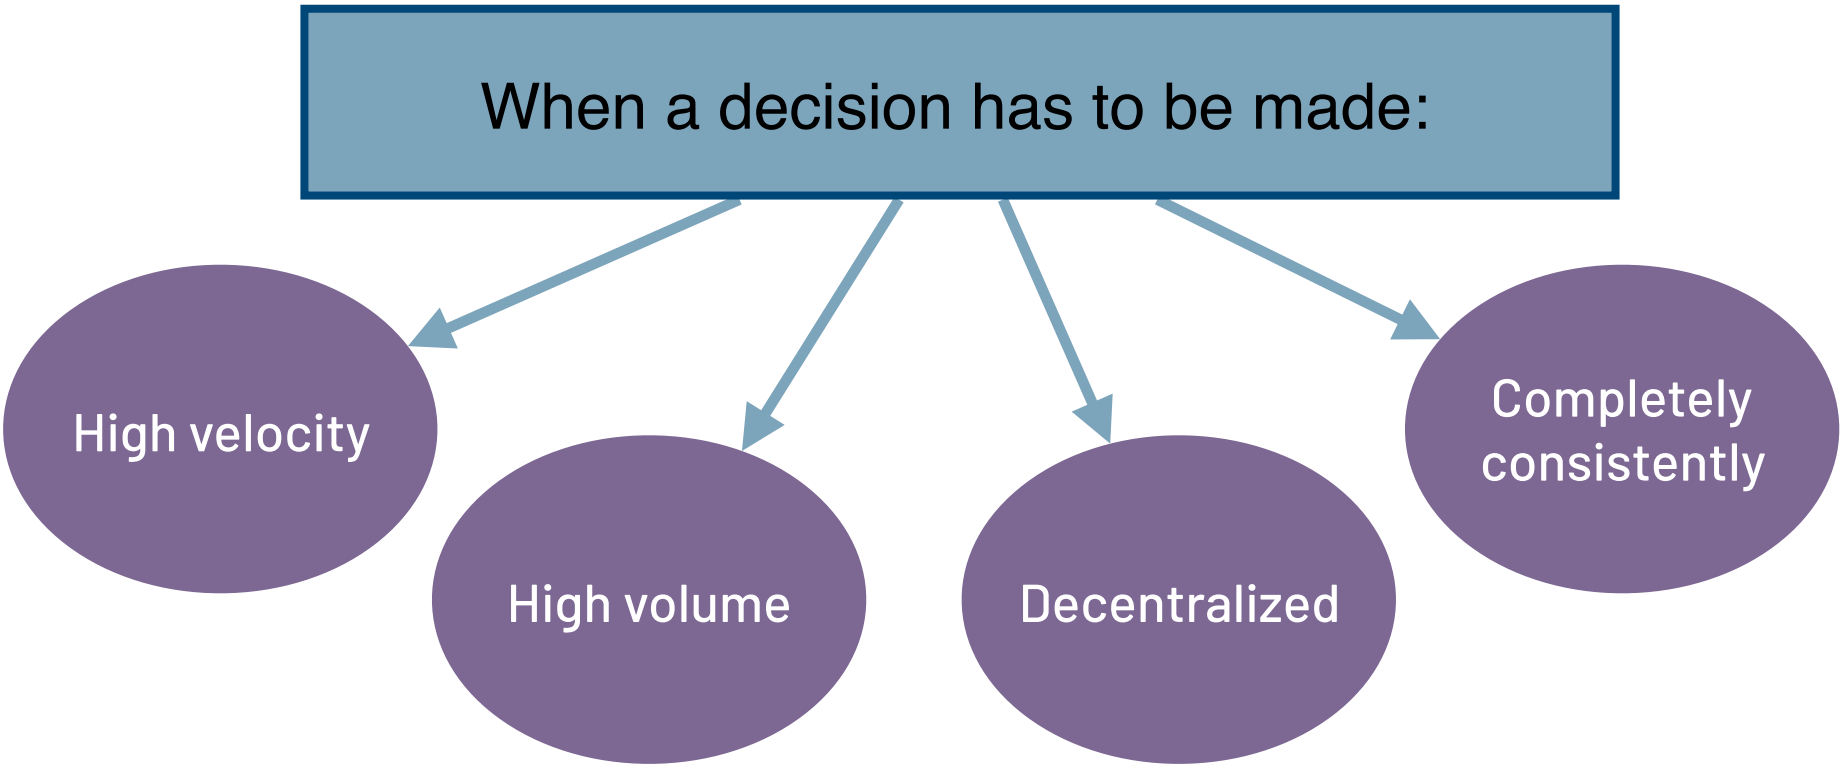

Successful Applications of Supervised Learning

"**Two-thirds** of the opportunities to use AI are in improving the performance of **existing analytics** use cases."

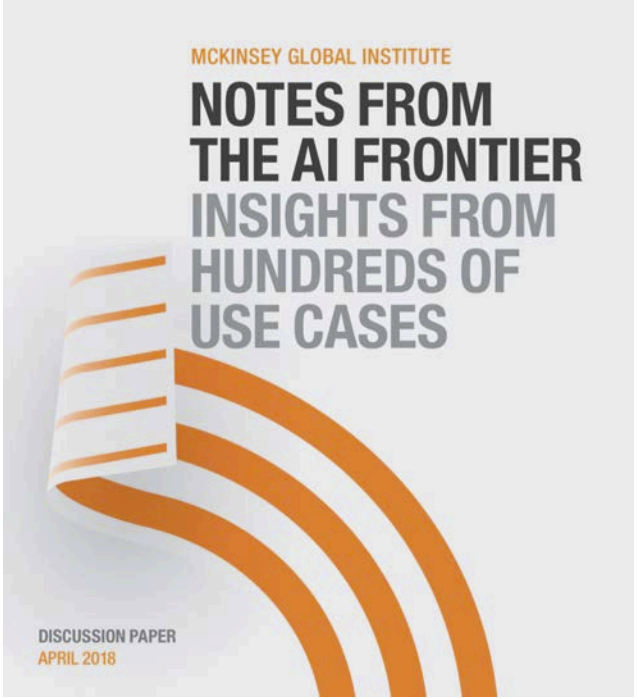

Innovative Applications and Amara's Law

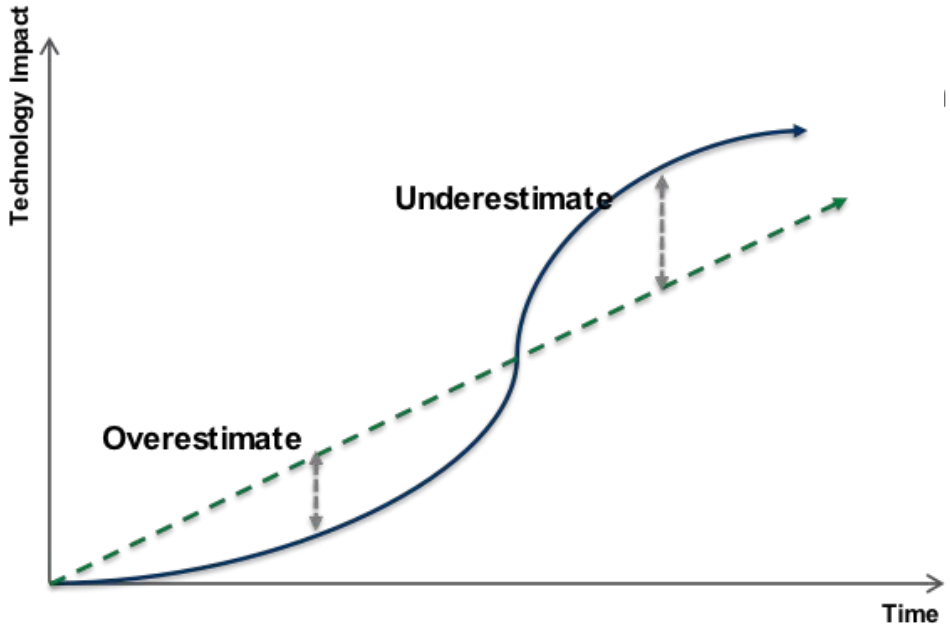

We tend to **overestimate** the effect of a technology in the **short** run and **underestimate** the effect in the **long** run.

https://medium.com/@QwQiao/technical-analysis-and-amaras-law-d456d2f356e7

Project Components

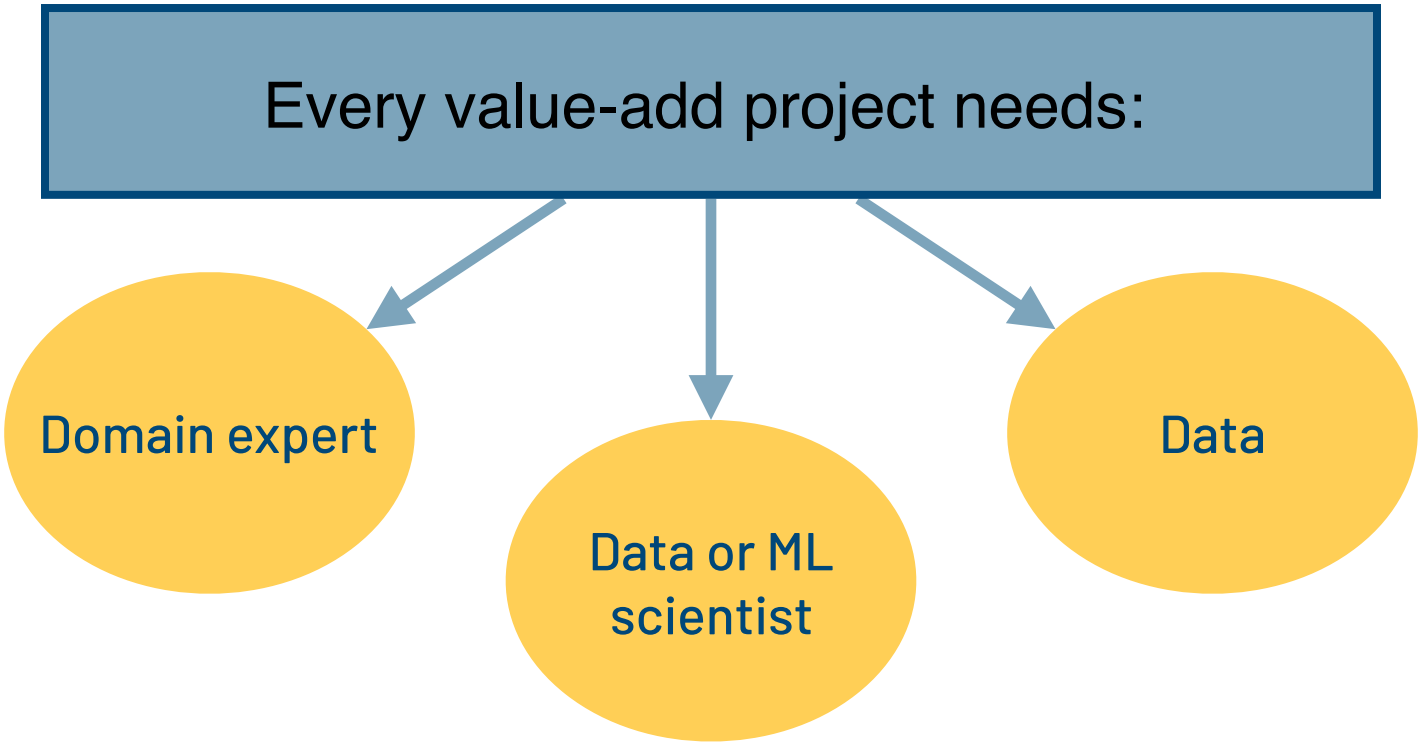

Where do you find projects?

ML project: idealized

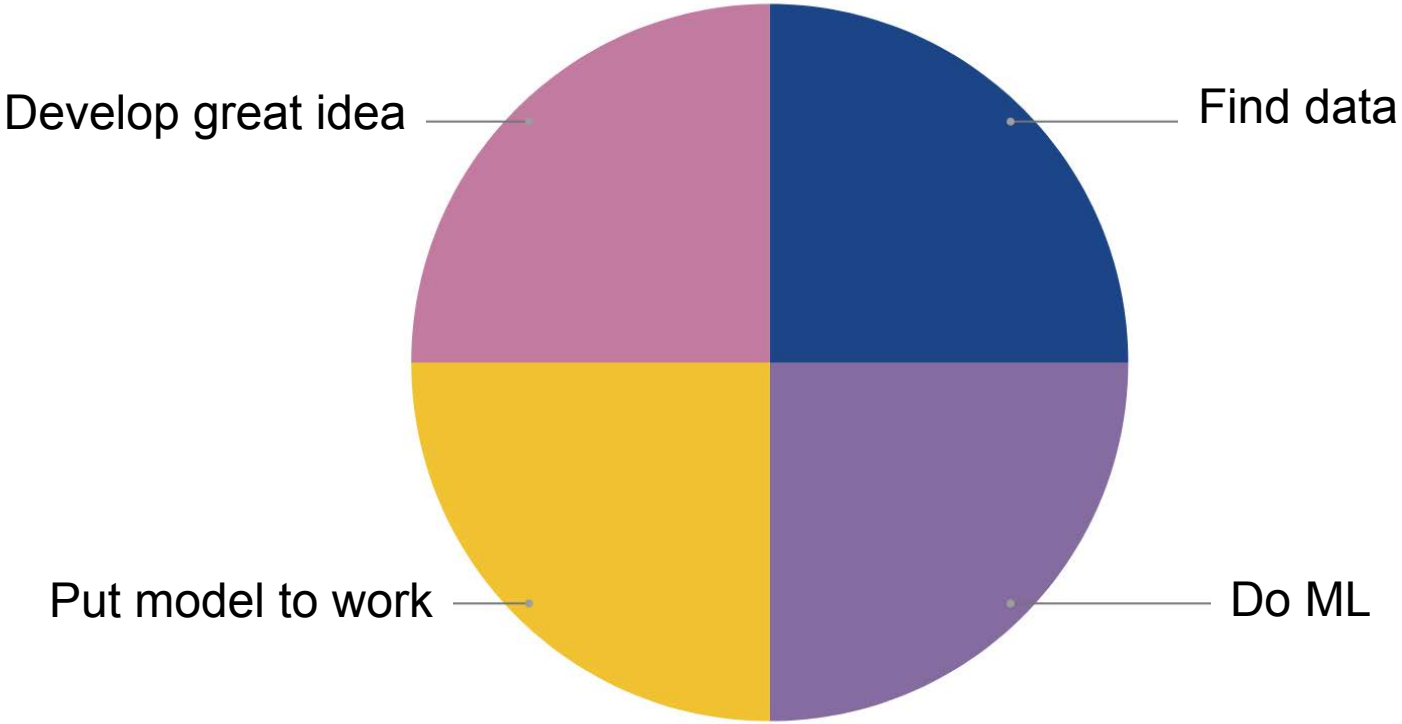

ML projects: reality

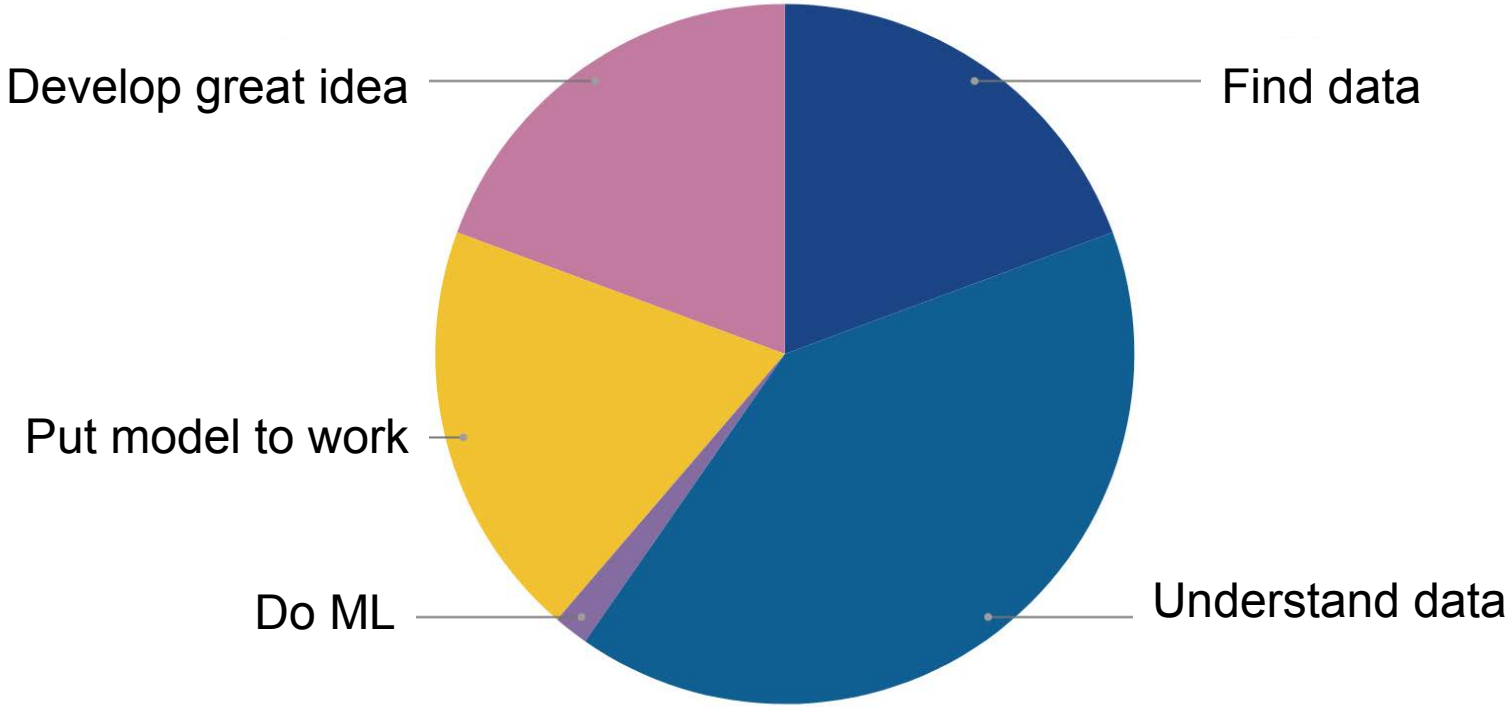

Case Study

Grocery store competition

kaggle Search Competitions Datasets Kernels Discussion Courses ••• Sign In

Overview Data Kernels Discussion Leaderboard Rules Join Competition

Overview

Description
Evaluation
Prizes
Timeline

Brick-and-mortar grocery stores are always in a delicate dance with purchasing and sales forecasting. Predict a little over, and grocers are stuck with overstocked, perishable goods. Guess a little under, and popular items quickly sell out, leaving money on the table and customers fuming.

The problem becomes more complex as retailers add new locations with unique needs, new products, ever transitioning seasonal tastes, and unpredictable product marketing. Corporación Favorita, a large Ecuadorian-based grocery retailer, knows this all too well. They operate hundreds of supermarkets, with over 200,000 different products on their shelves.

Corporación Favorita has challenged the Kaggle community to build a model that more accurately forecasts product sales. They currently rely on subjective forecasting methods with very little data to back them up and very little automation to execute plans. They're excited to see how machine learning could better ensure they please customers by having just enough of the right products at the right time.

How would you (now) describe machine learning?

Other Amii Courses

ML Foundations

June 19-21, 2019

ML Technician Program

September 24, 2019

(8 month program)

Thank you!

Questions?

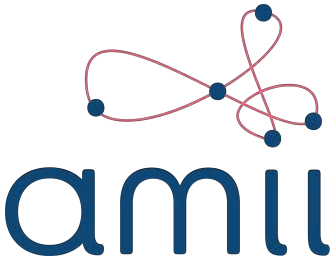

hello@amii.ca
www.amii.ca

Main takeaways

What is ML?

AI, ML, Classic AI, QuAMs
Supervised, Unsupervised,
Reinforcement Learning

What can it do?

Prediction (number of sales)
Classification (on-time delivery)
Association (market basket analysis)
Action Control (robo warehouse)
Planning (route finding)

How to succeed?

Data does not equal knowledge
Garbage in, garbage out
Automate when beneficial
Projects need expertise

Case Study

Kaggle competitions
Feature engineering
Data stratification/number of models
Optimization criteria