Lady Ada Lovelace

AI for Decision Making

Phanish Puranam, INSEAD

Phanish Puranam

The Roland Berger Chaired Professor of Strategy and Organisation Design

- PhD in Management from the Wharton School of the University of Pennsylvania
- Served as the Academic Director for the PhD Program at both London Business School and INSEAD
- Expertise: Organization Design & Strategy
- Current research: Organizations & Algorithms, Remote Collaboration

What does it take to master AI?

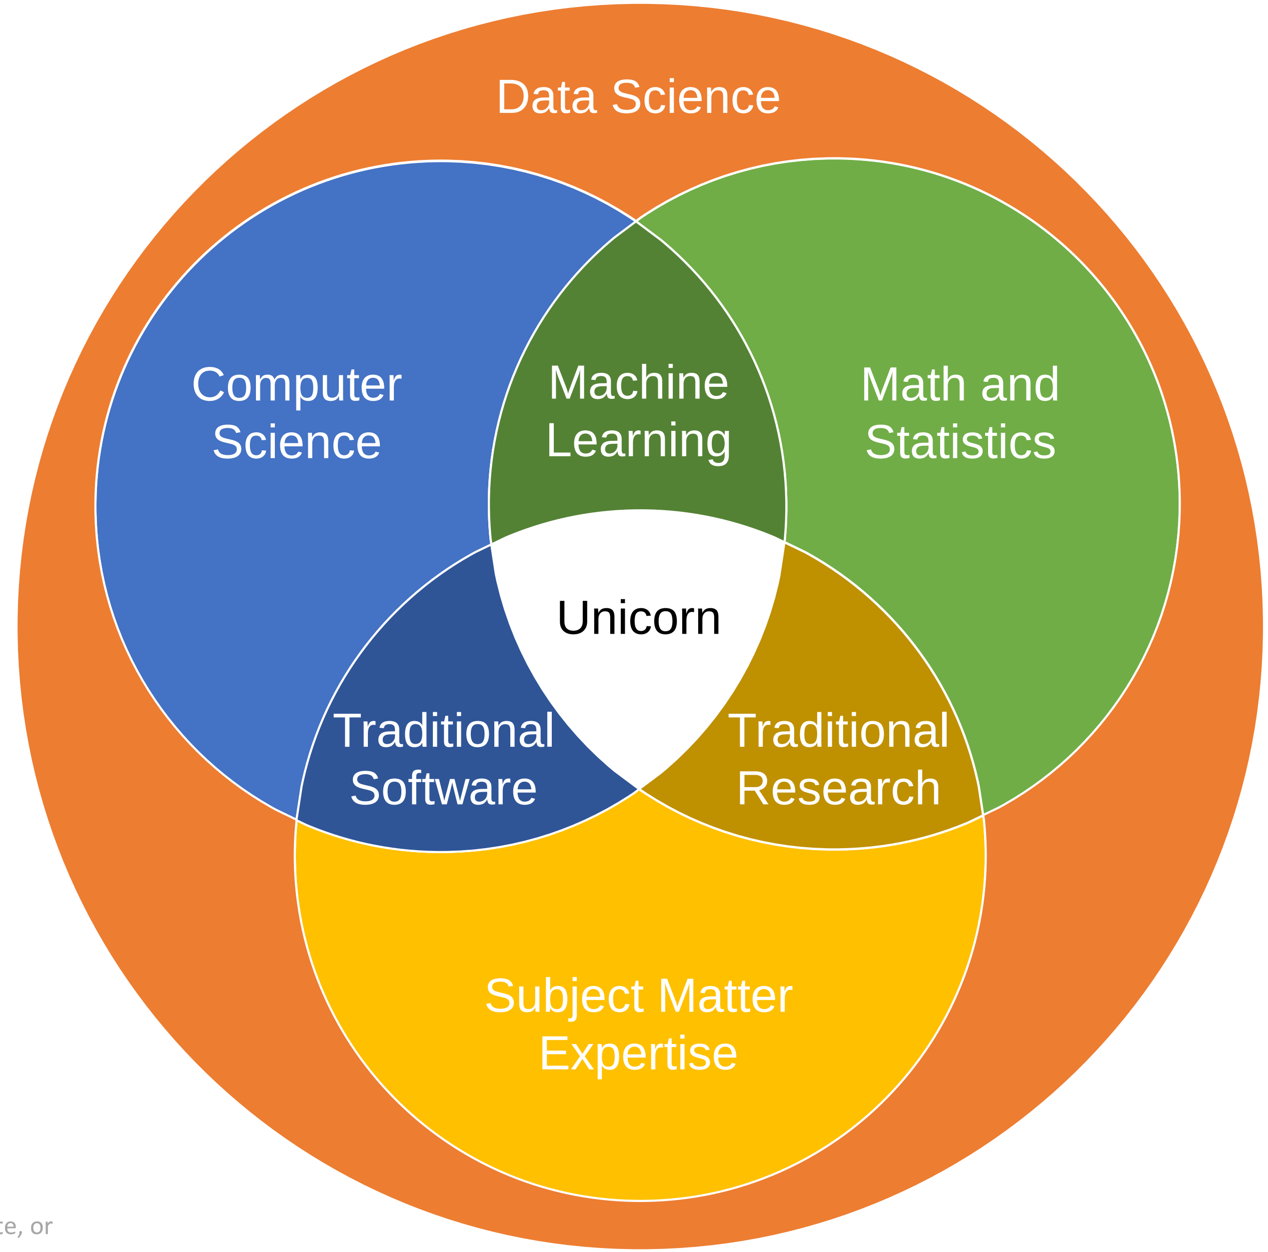

What’s the <u>**most**</u> impressive accomplishment of AI to date (in ***your*** view)?

- Please enter into our Shared Google Doc- do check for redundancy with what others are typing!
- 3 minutes

What’s the <u>**most**</u> impressive accomplishment of AI to date (in ***your*** view)?


PYRAMID
PREDICTIVE ANALYTICS
BUSINESS
BANKING AND FINANCE
RETAIL
TELECOMMUNICATIONS
CREDIT SCORING


Automate
Workflow
Analyze
Monitor

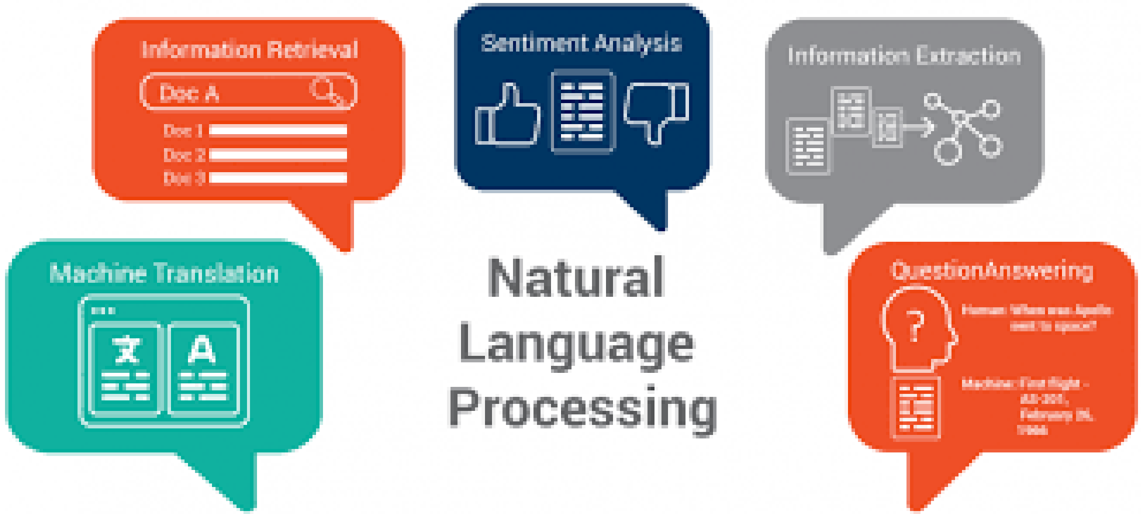
Information Retrieval
Doc A
Doc 1
Doc 2
Doc 3
Sentiment Analysis
Information Extraction
Machine Translation
Natural Language Processing
QuestionAnswering

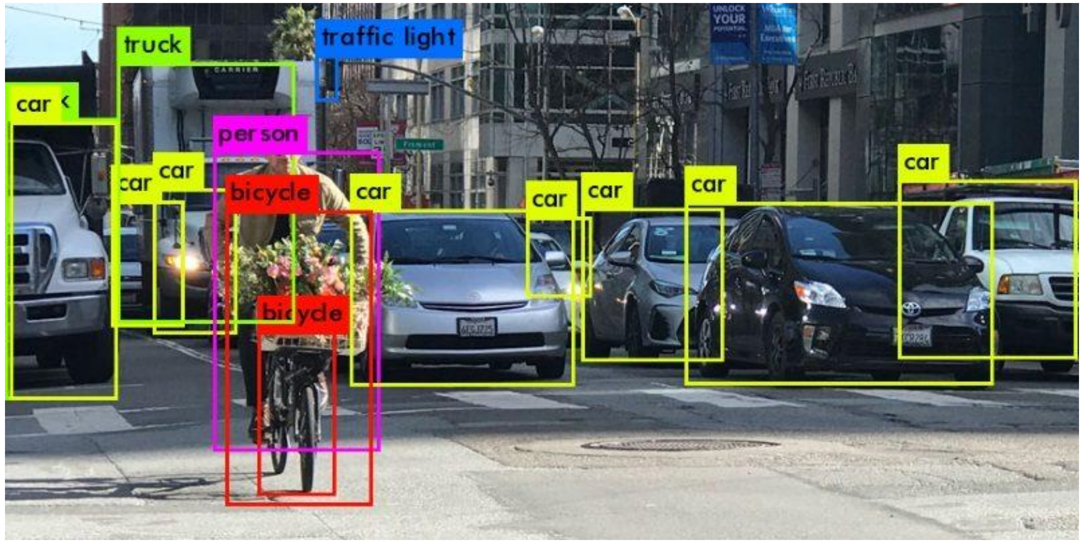
truck
traffic light
car
person
car
car
bicycle
car
car
car
car
car
bicycle

Can AI be creative?

Figure 5: 1024 × 1024 images generated using the CELEBA-HQ dataset. See Appendix F for a larger set of results, and the accompanying video for latent space interpolations.

Progressive growing of GANs

1959

1997

2016

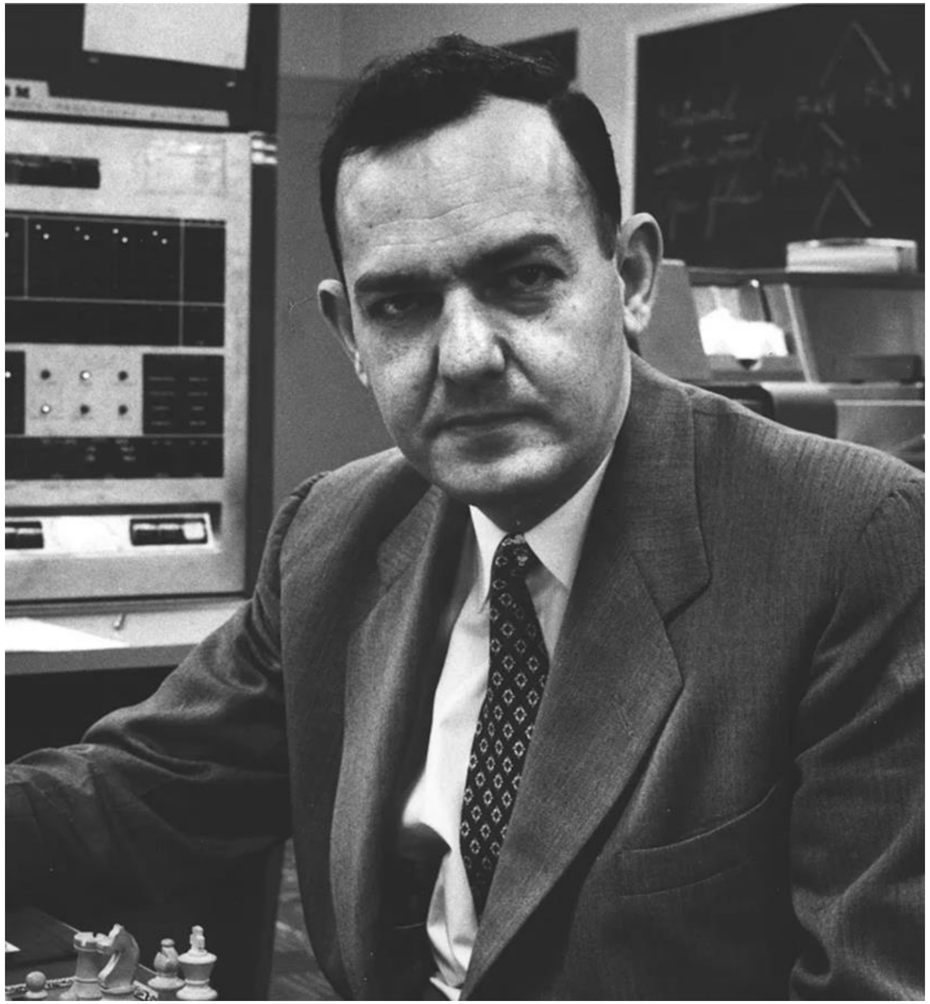

Herbert Simon (joint work with W.G. Chase, 1973)

HUMAN EXPERTISE= INTELLIGENT BEHAVIOR IN A CONTEXT= PATTERN RECOGNITION !

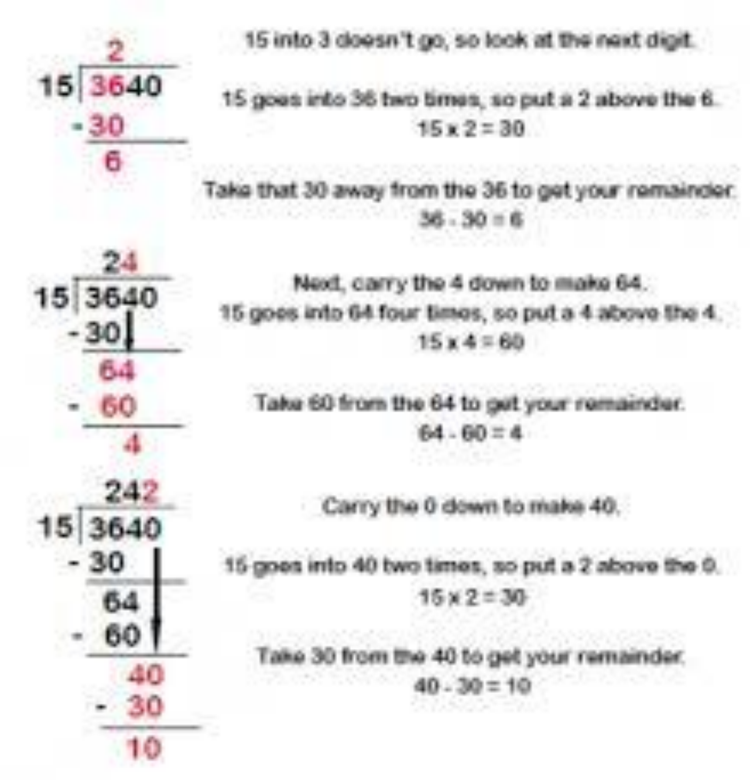

OLD AI: EXECUTES RULES

NEW AI: LEARNS PATTERNS FROM DATA

Name: ______________________

Complete the Pattern

Look at each pattern below and then use the box on the right to complete the pattern.

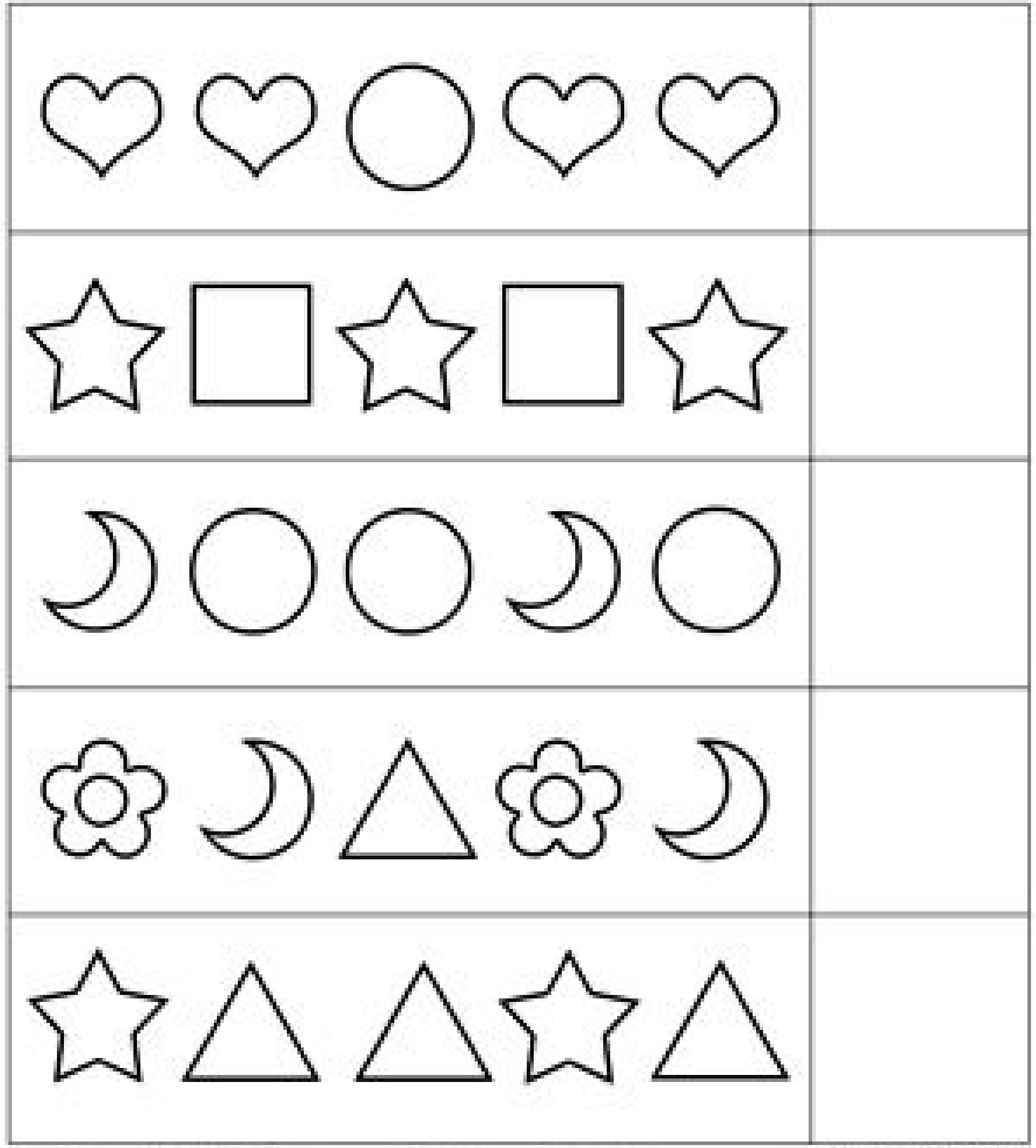

Pattern recognition & application

- Learning to recognize and apply patterns is the heart of intelligence
- That's what Machine Learning does
- Lots of data + processing power makes it easier to recognize (even complex) patterns
- A pattern= a "function" – which associates inputs to *unique* outputs.
- Pattern recognition in humans = function approximation in algorithms

Machine learning ⊆ artificial intelligence

ARTIFICIAL INTELLIGENCE

Design an intelligent agent that perceives its environment and makes decisions to maximize chances of achieving its goal.
Subfields: vision, robotics, machine learning, natural language processing, planning, …

MACHINE LEARNING

Gives "computers the ability to learn without being explicitly programmed" (Arthur Samuel, 1959)

SUPERVISED LEARNING

Classification, regression

UNSUPERVISED LEARNING

Clustering, dimensionality reduction, recommendation

REINFORCEMENT LEARNING

Reward maximization

A STORY OF DOGS AND CATS

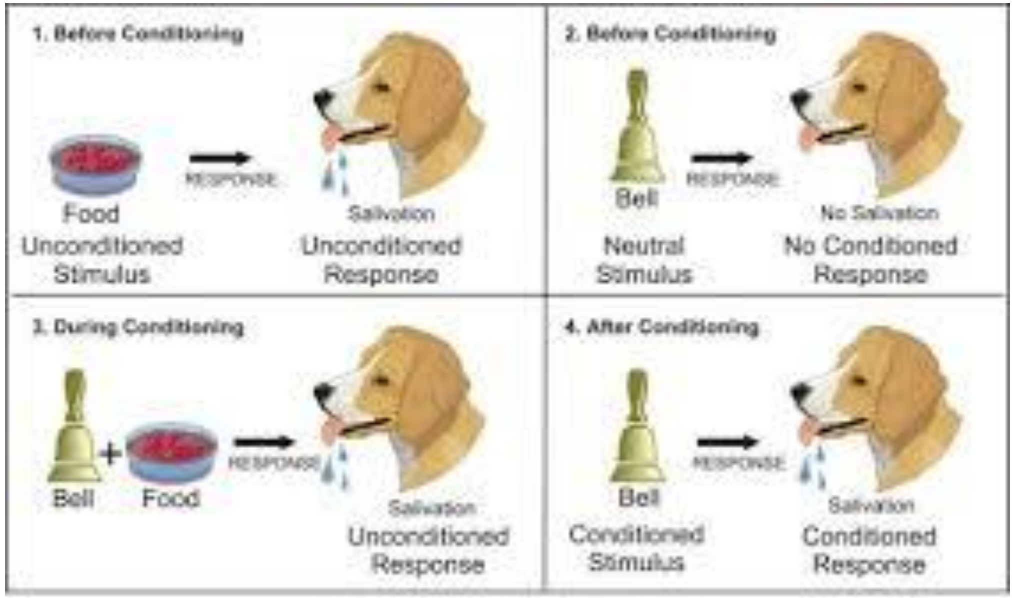

Classical Conditioning

Figure 18.2 Thorndike's puzzle box

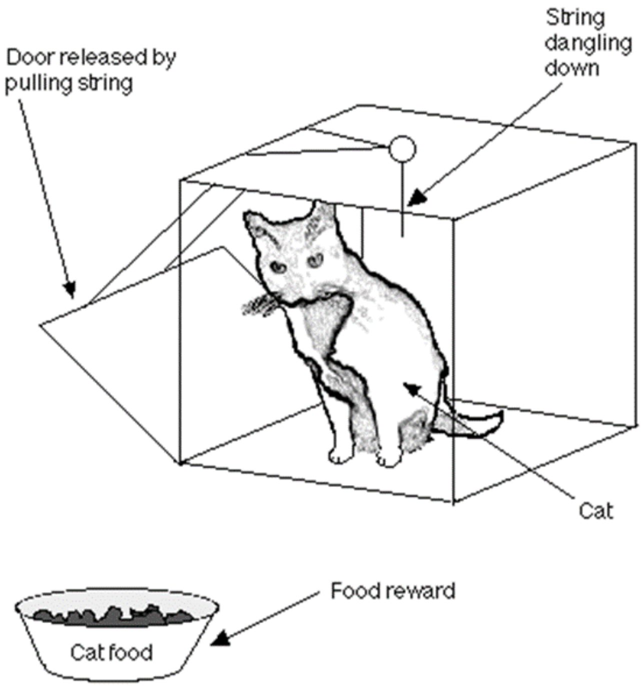

Supervised & Unsupervised Learning

Reinforcement Learning

ML algorithms work through pattern recognition & completion

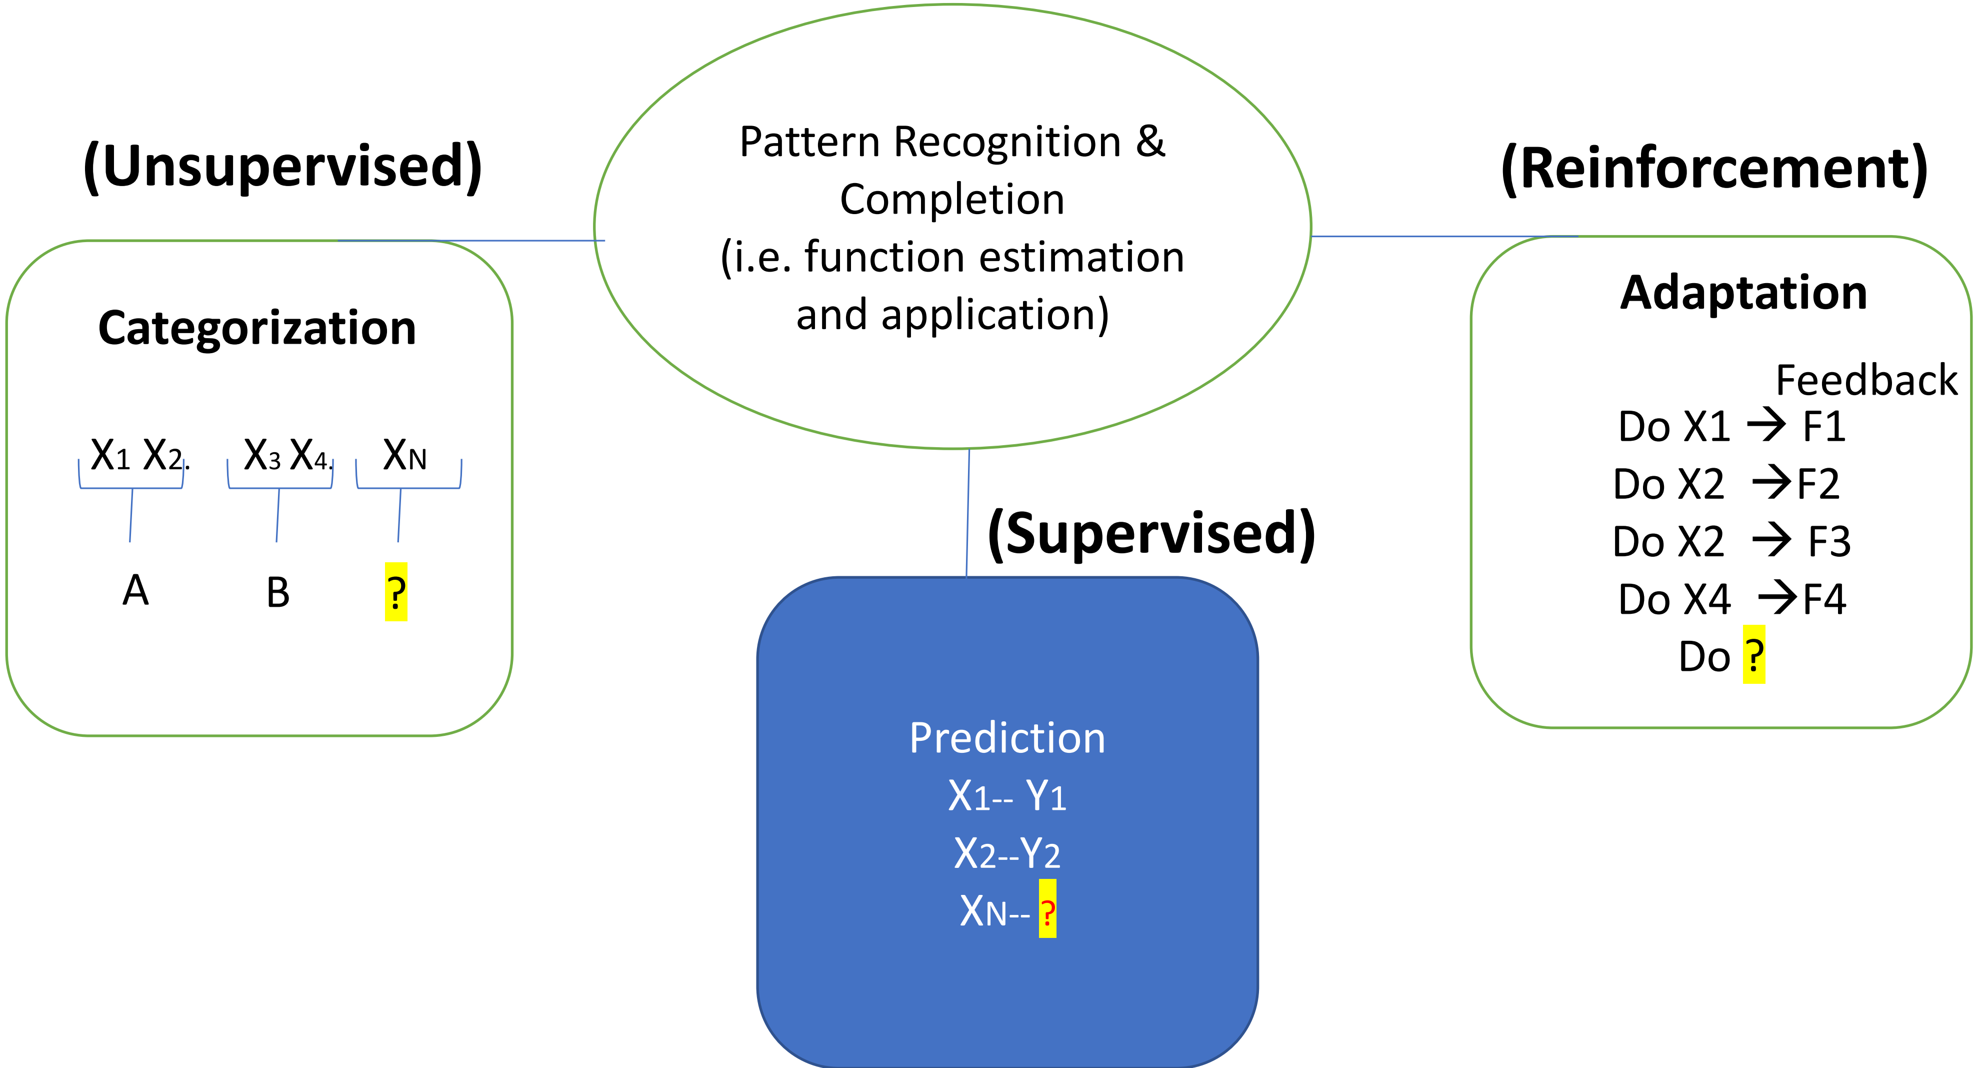

Patterns → Predictions → Decisions

- All decisions require prediction

- Decisions can be “standalone” vs. “embedded”

AI ~ ML~ Pattern Detection: Standard use cases today

Transaction features ---- Fraud

Customer attributes---- Churn likelihood

Inkdots ---- Words

AI ~ ML~ Pattern Detection: Standard use cases today

Pixel ----- Image

Sound ------ Words in speech

https://www.youtube.com/watch?v=D5VN56jQMWM&t=127s

State of play ----- Optimal next move in game

https://www.eff.org/ai/metrics#Abstract-Strategy-Games

AI ~ ML~ Pattern Detection : Emerging Use cases

Employee attributes----Exit likelihood

CV attributes ---- Good hire

Project characteristics ---- Good investment

AI ~ ML~ Pattern Detection : Emerging Use cases

Text and word use ----- Cultural values

Company characteristics----- Strategic moves

Hans Moravec's landscape, in Tegmark (2018)

Dominance or Co-existence?

37

78


8 - 15 March 2016
ALPHAGO
01:58:37
AlphaGo
Lee Sedol

HOW TO KEEP UPTODATE ON WHAT AI IS DOING (WELL)

https://www.eff.org/ai/metrics

Levels of sophistication in using Data in Business

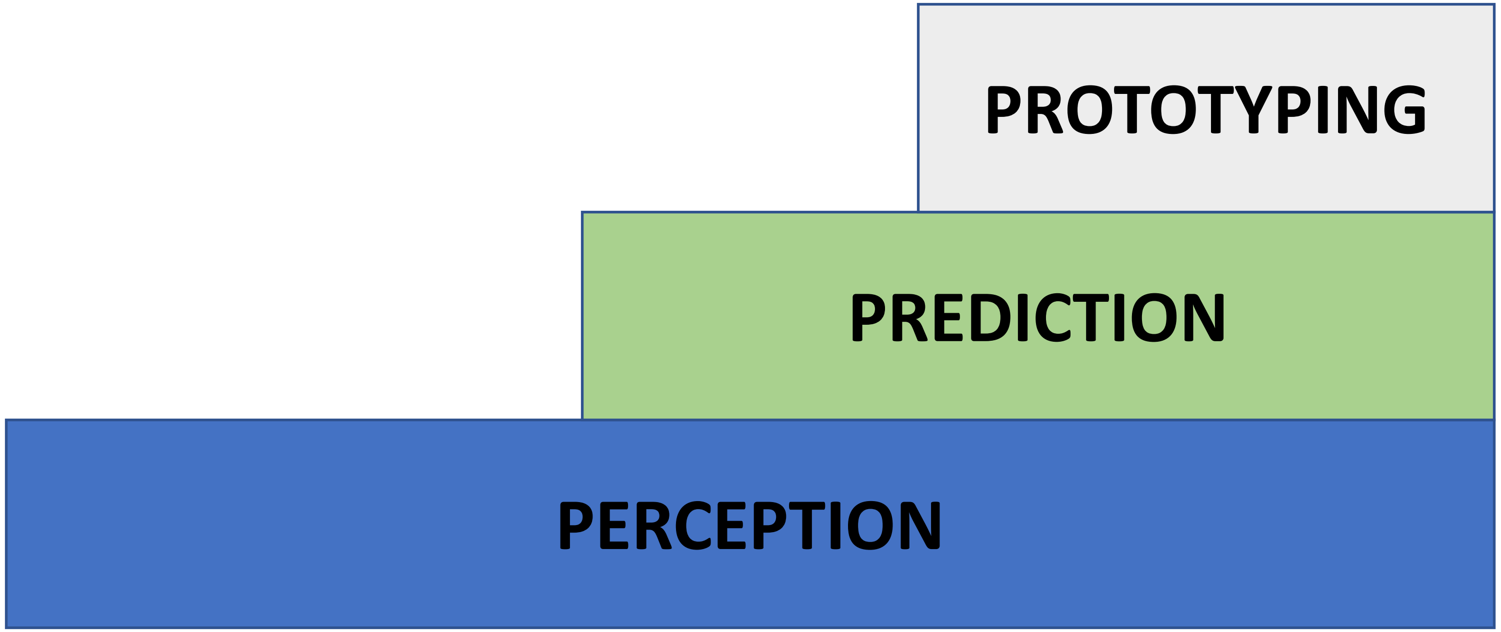

Data
Continuous Outcome

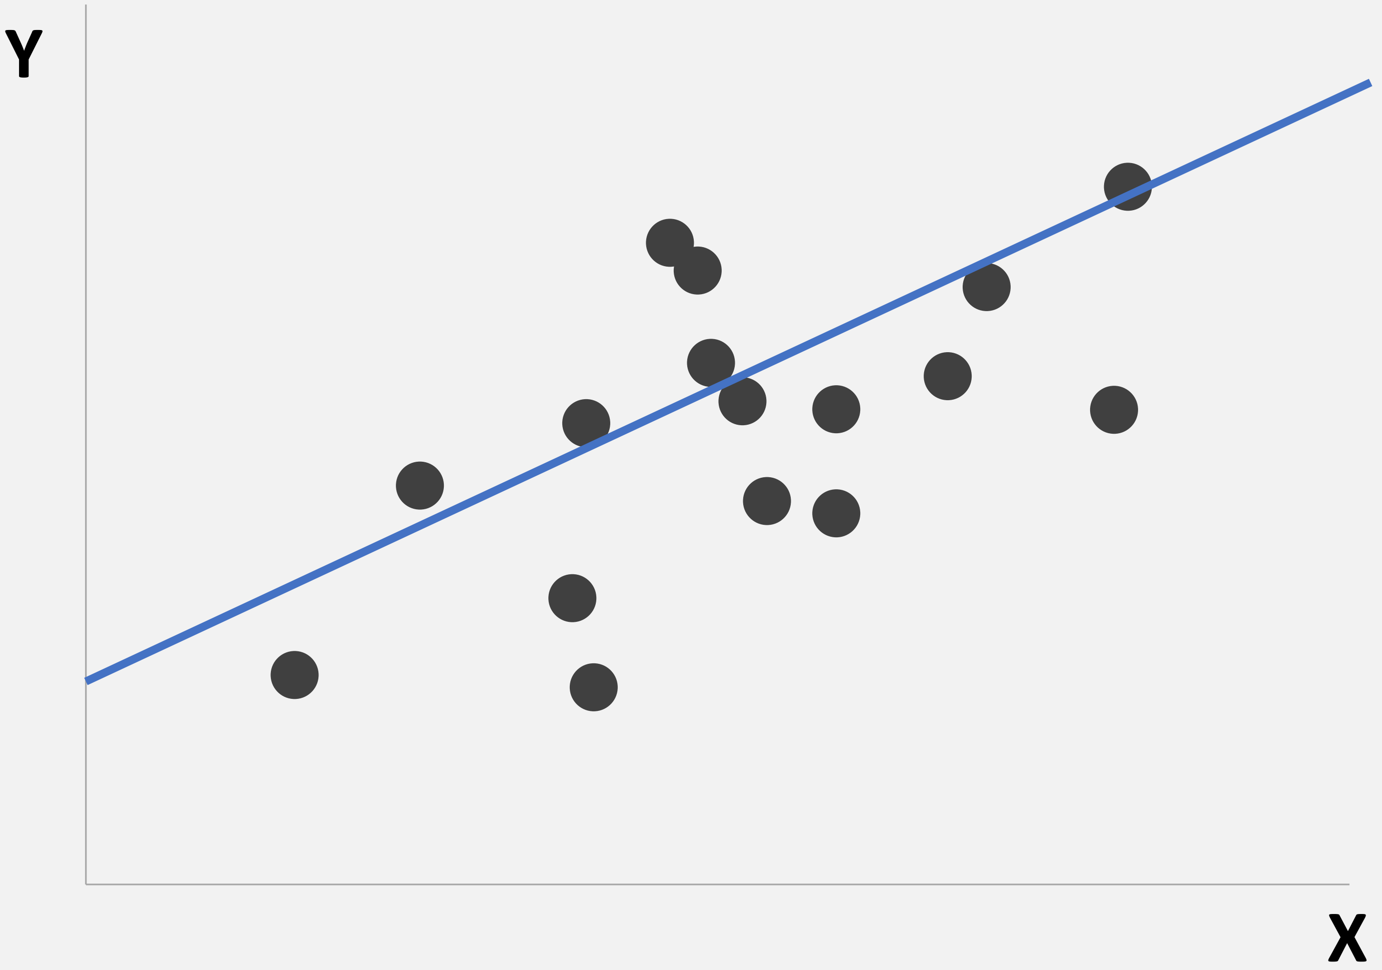

PERCEPTION vs. PREDICTION

- Perception through Hypothesis testing:
- Does this association exist in the data?
  - We have a hunch (hypothesis) about an explanation
  - **Null Hypothesis Significant Testing** tells us the probability (across many samplings) of observing the association we see merely by chance, even if the true association in the population is zero.
  - **p-value**

- Prediction through Data mining:
- What associations exist in the data?
  - We have no or minimal hunches, we “let the data speak” to make a prediction
  - **Data mining** involves algorithms that hunt for useful associations in the data that can be summarized in a model.
  - **Predictive accuracy**

Prediction What's going to happen?

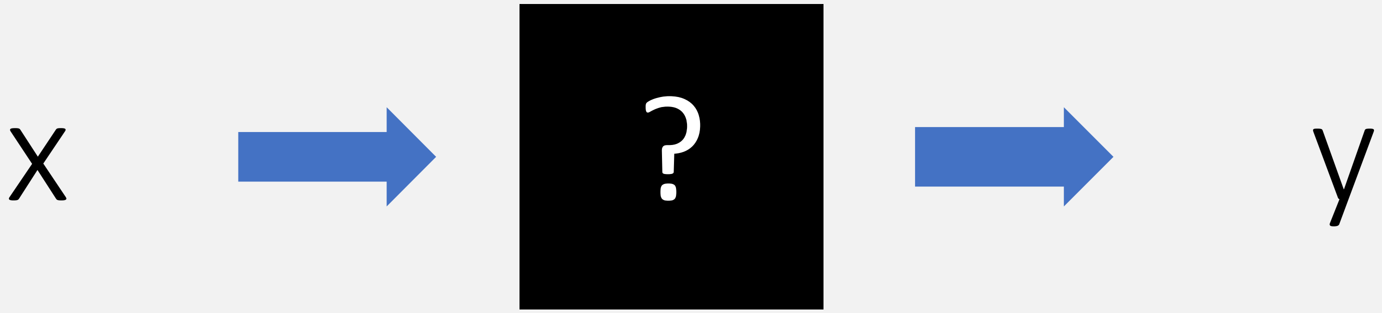

- Different algorithms
  - Tree induction
  - LASSO
  - Random forest
  - OLS
  - Logistic
  - Deep learning

- Common to all:
- **Use past data to predict**
- **future outcomes**

How can more advanced ML algorithms improve on what we just did?

Past Data
Continuous Outcome

Key issue: build models that fit current data well but also predict future data well

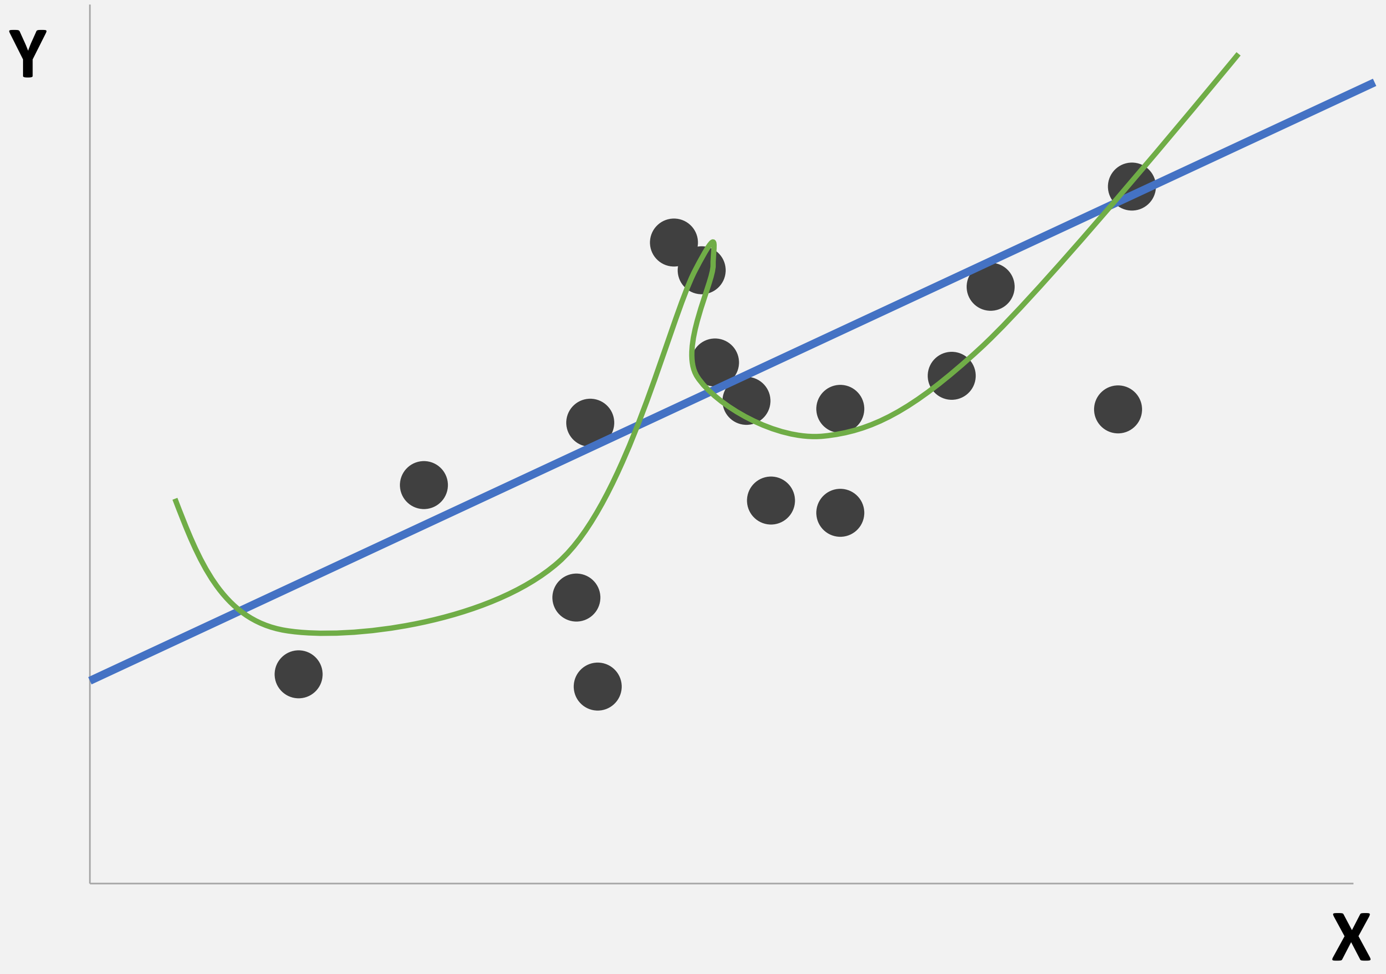

Overfitting + Generalization

- **Problem:**
- Every sample is truth + error. The model may learn too much about the error in a particular sample by adding too many parameters.

- **Solutions:**

1. **Penalizing predictive accuracy for model complexity**
   1. Regularization (constraint on parameters)
2. **Partition training data into training and test data**
   1. K-fold validation
3. **Both!**

Models with more predictive power are often more opaque.

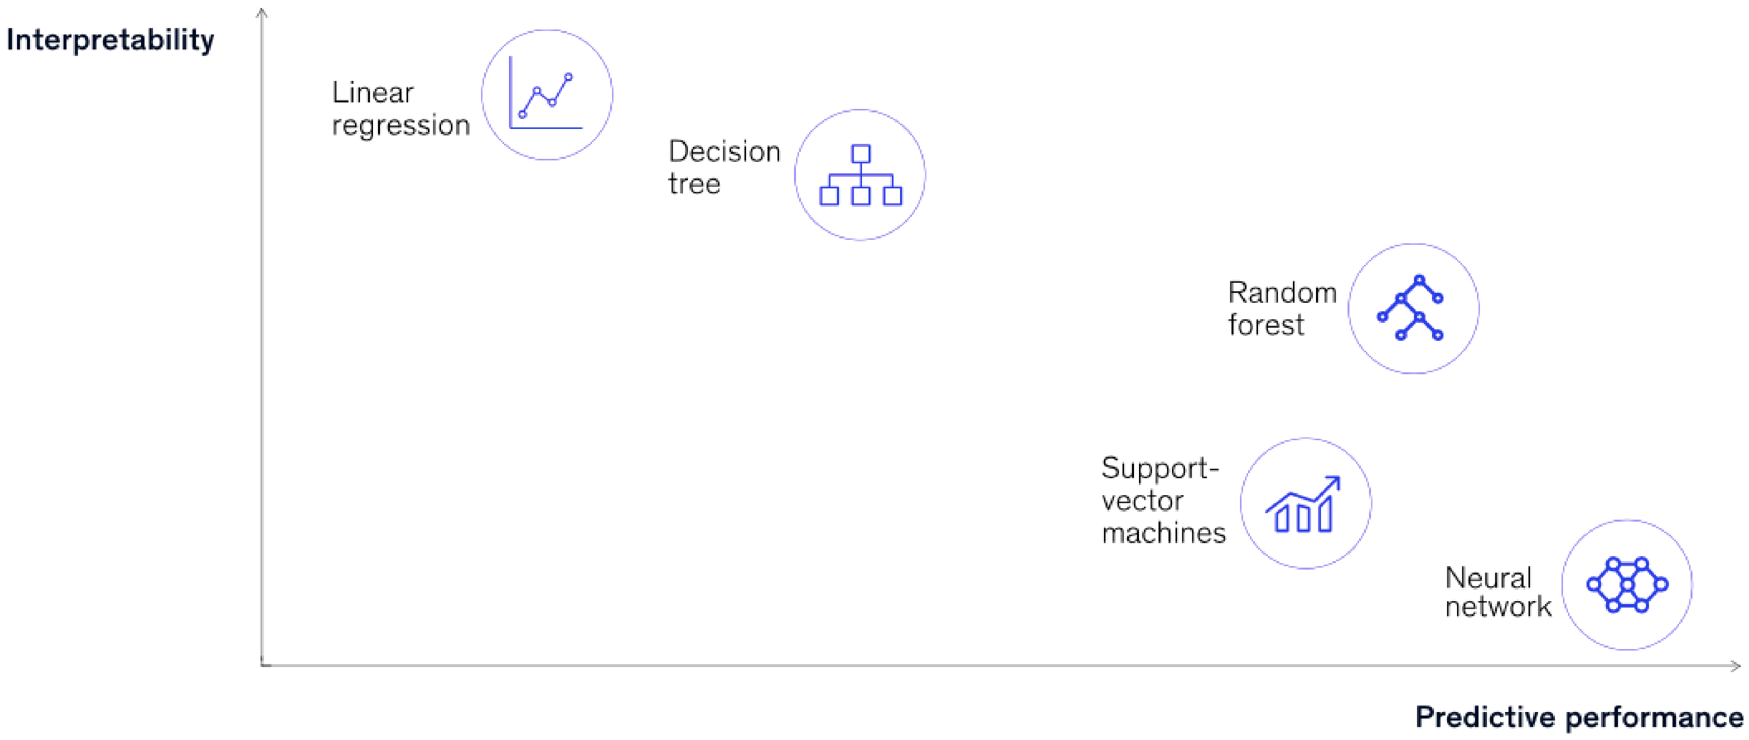

McKinsey & Company

Components of a prediction machine

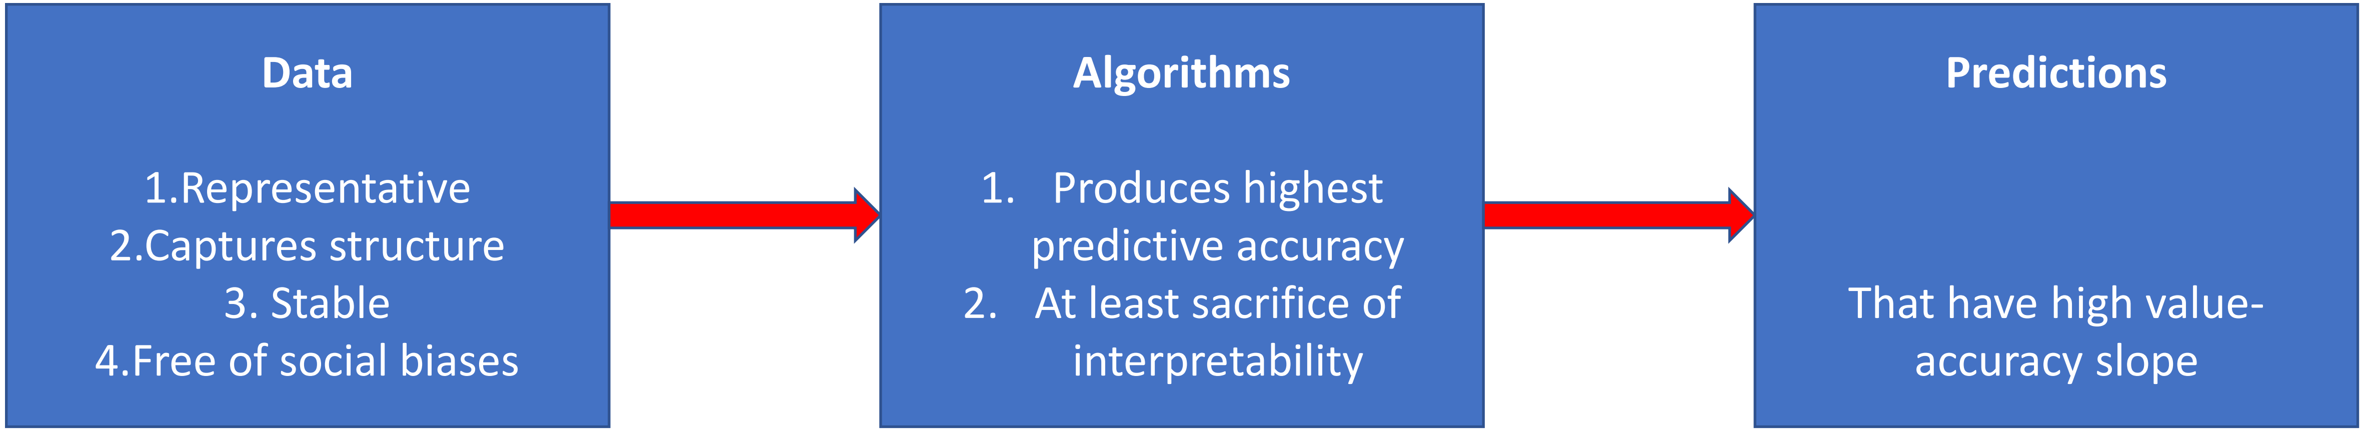

https://knowledge.insead.edu/blog/insead-blog/where-ai-can-help-your-business-and-where-it-cant-13136

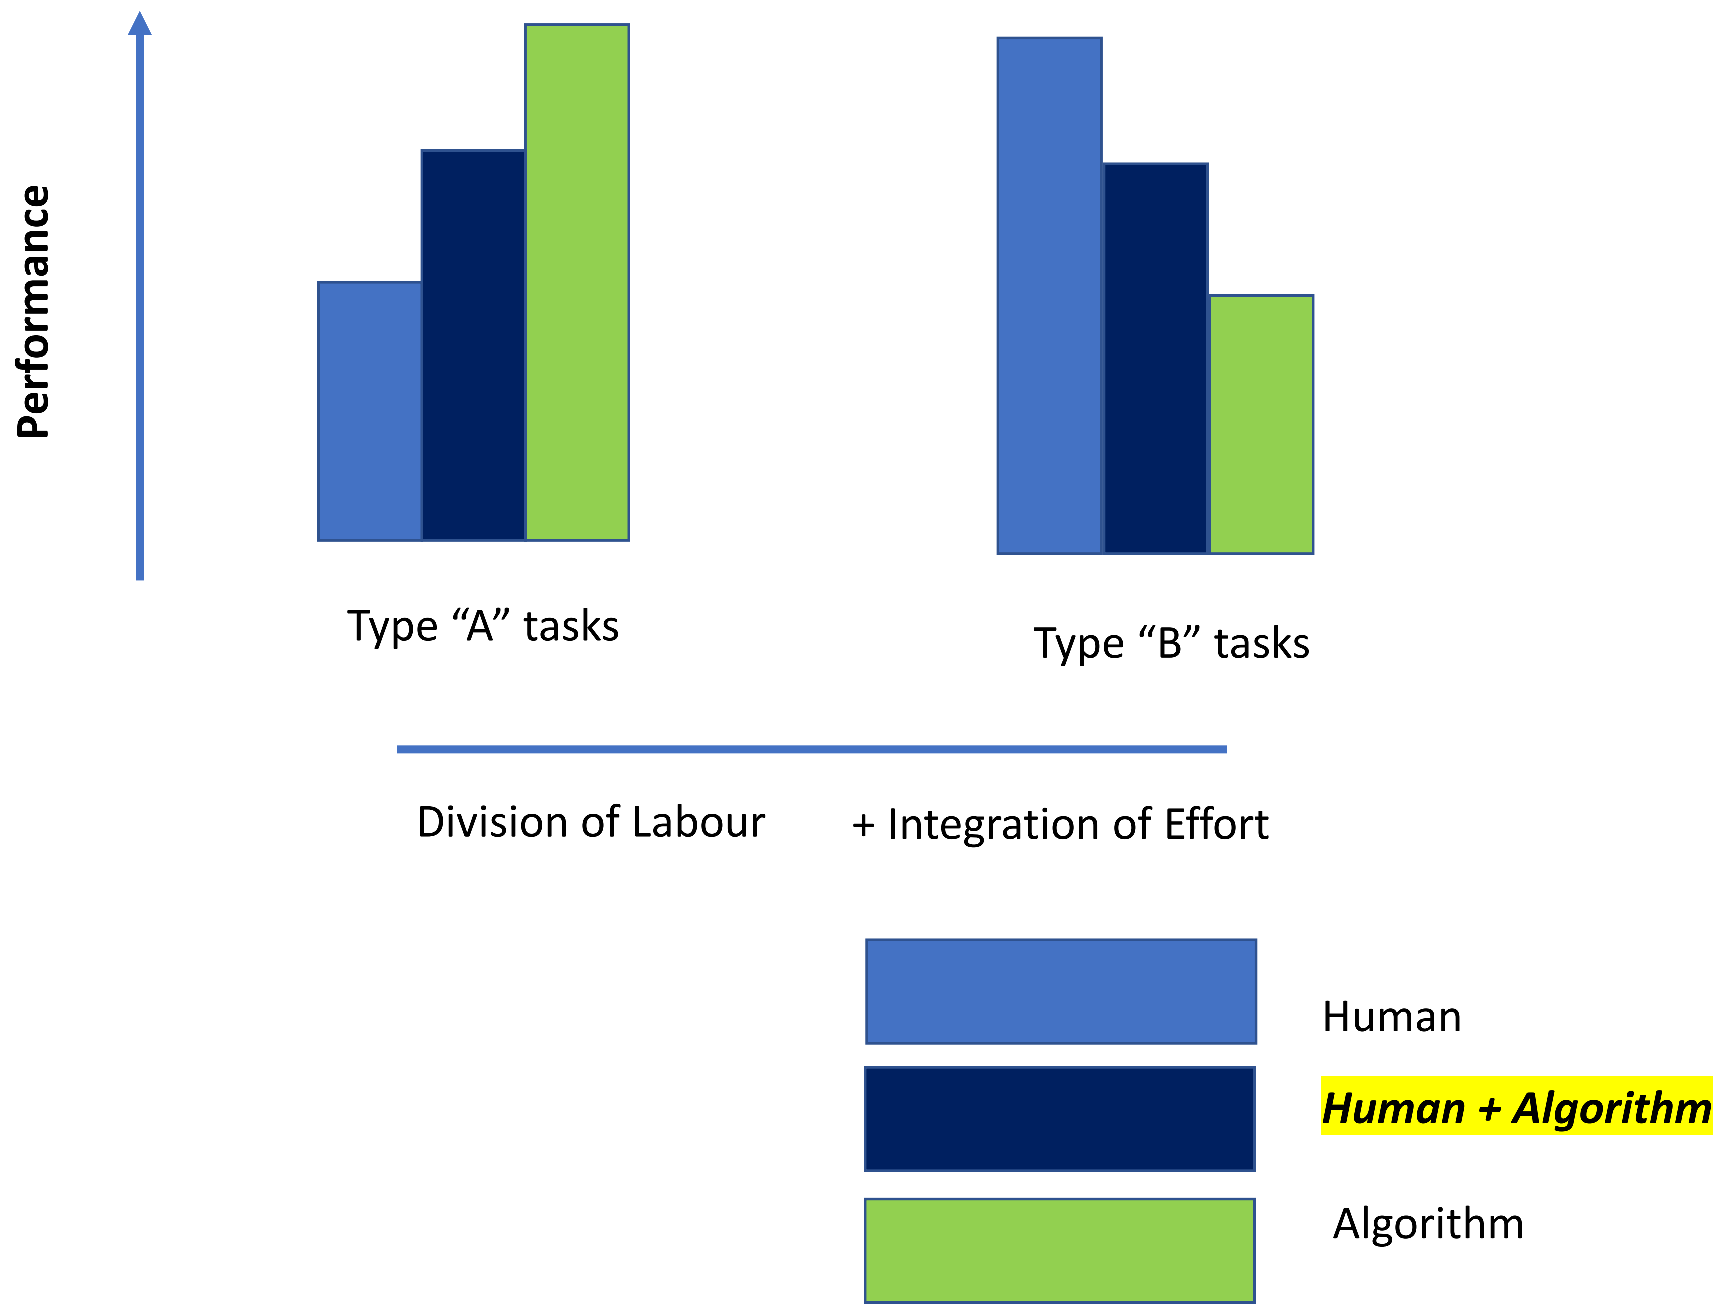
Performance
Type “A” tasks
Type “B” tasks
Division of Labour
+ Integration of Effort
Human
Human + Algorithm
Algorithm

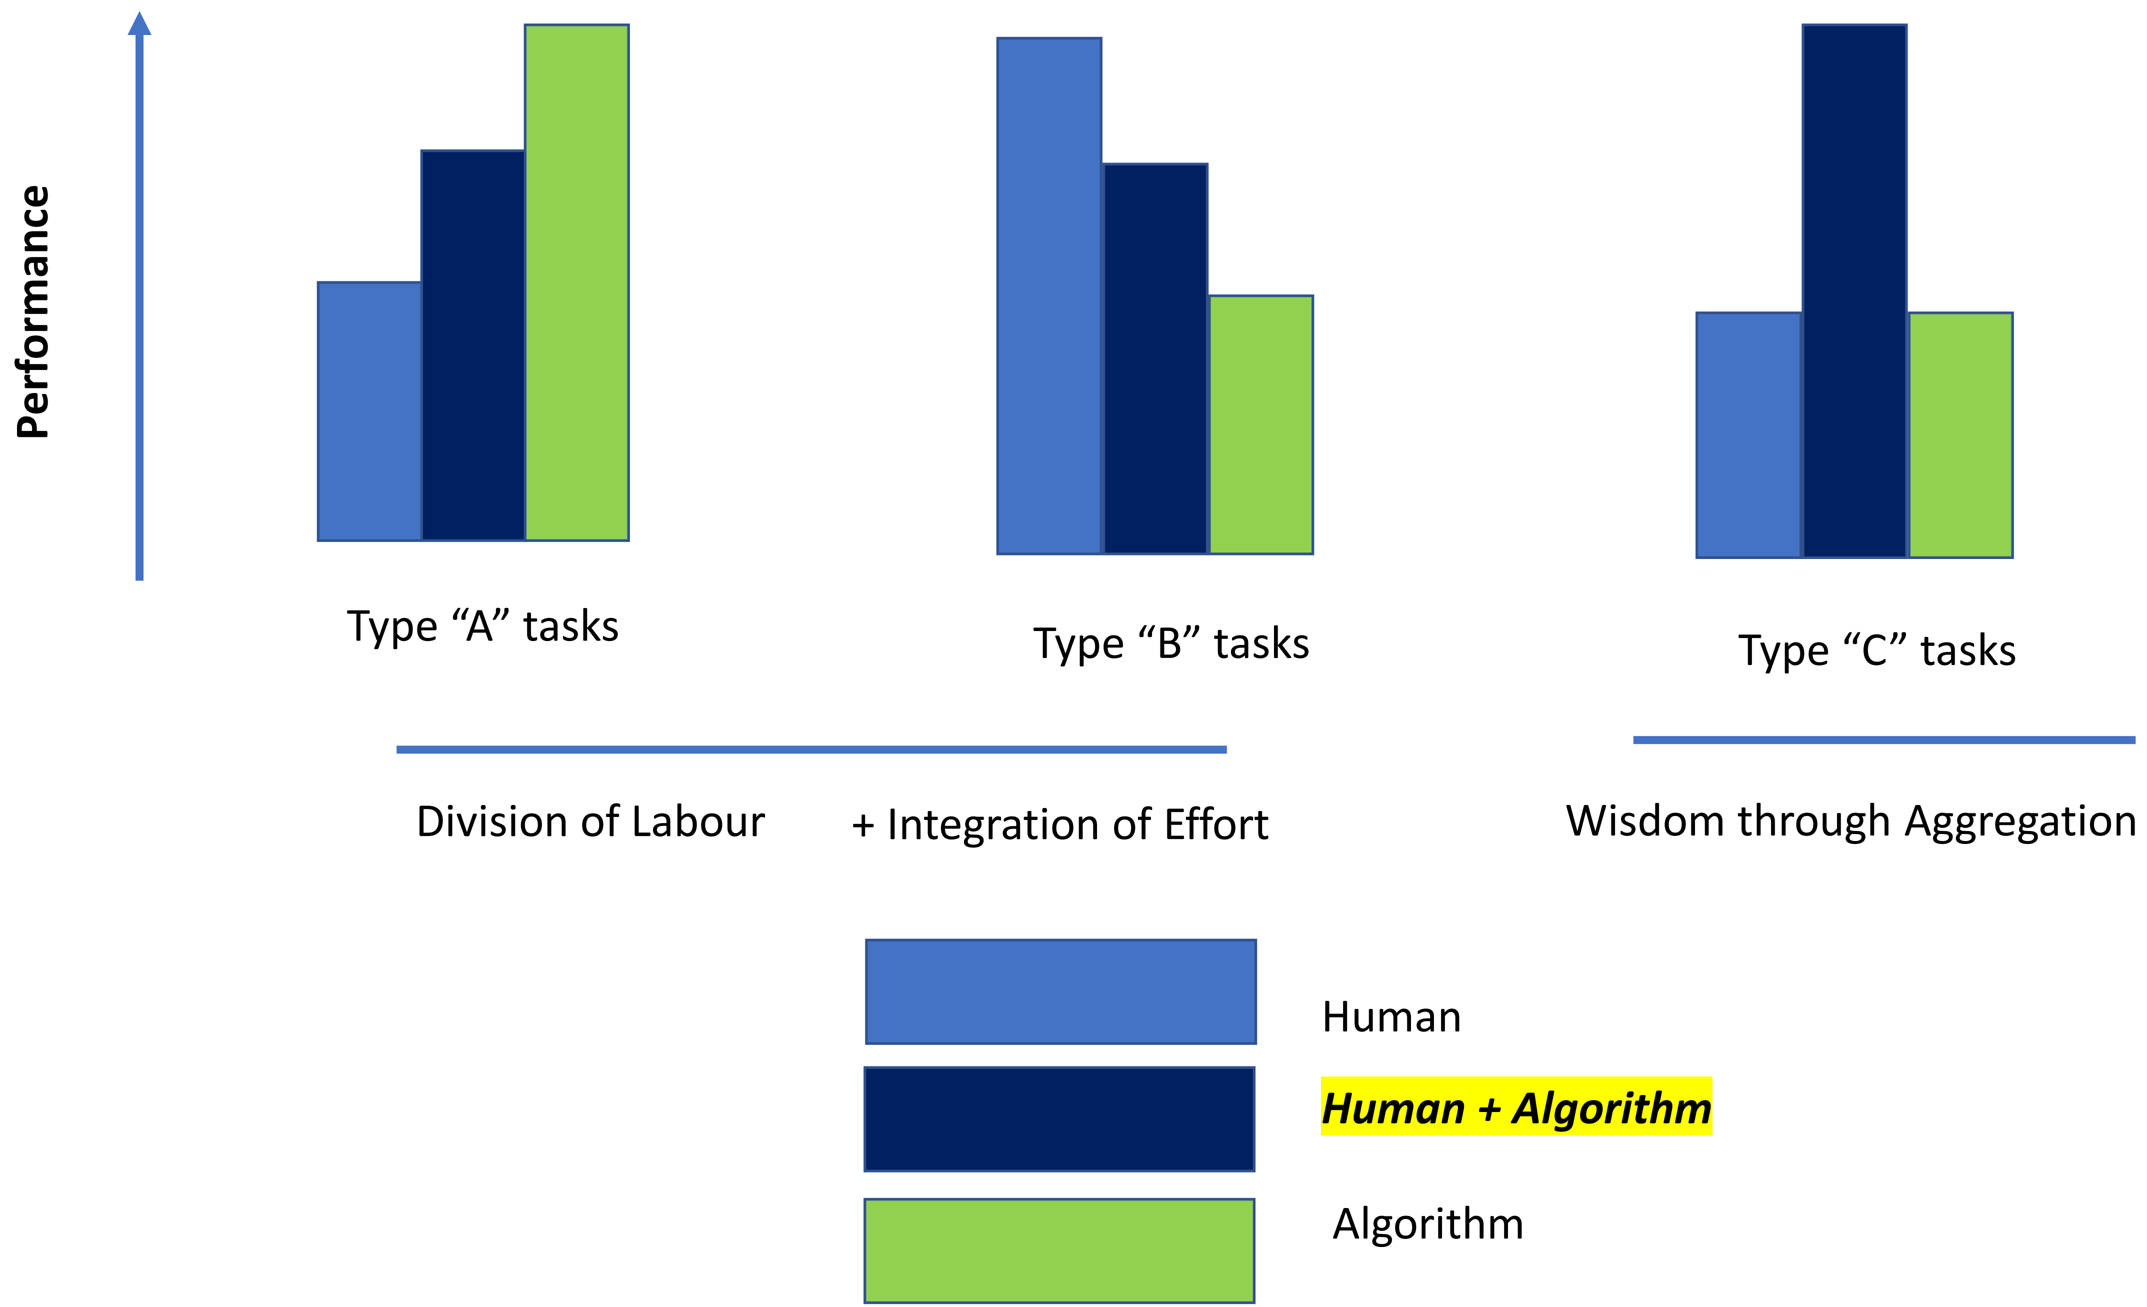
Performance
Type "A" tasks
Type "B" tasks
Type "C" tasks
Division of Labour
+ Integration of Effort
Wisdom through Aggregation
Human
Human + Algorithm
Algorithm