AI for Everyone: Demystifying the Basics of Artificial Intelligence

Yanou Ramon, Ph.D.

Business Contracts & Technology
University of Antwerp
6 May 2024

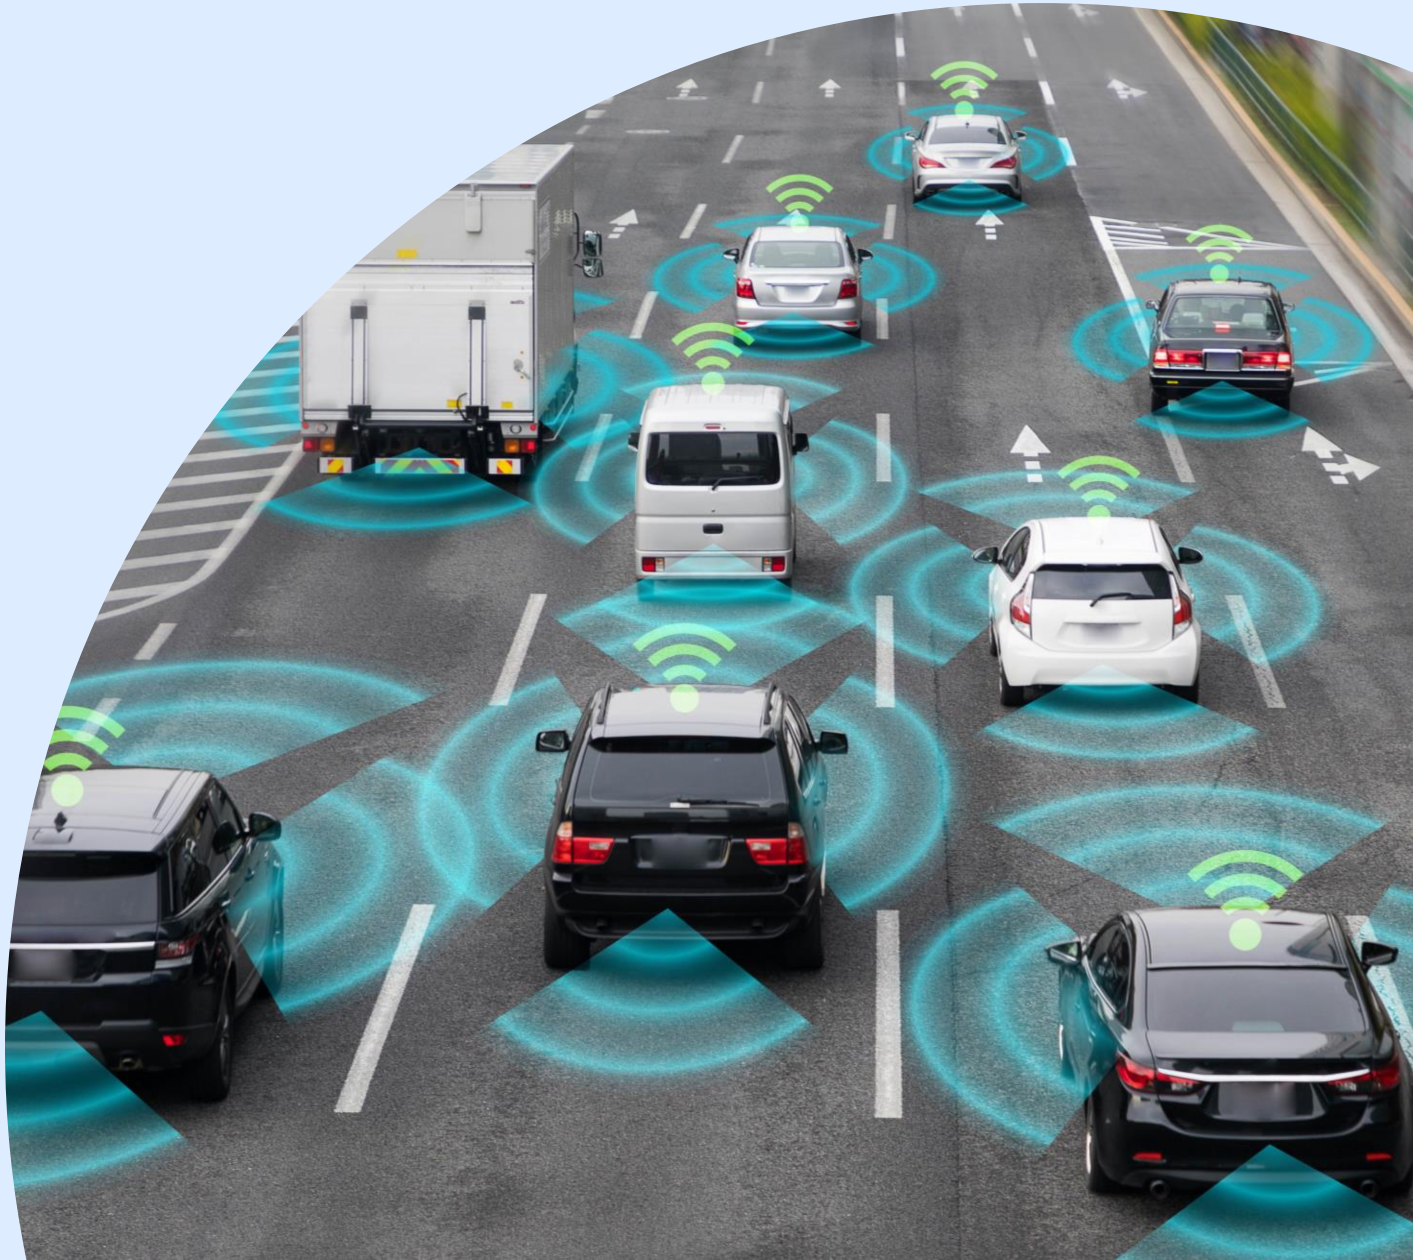

About me

2013-18: MA in Business Engineering (University of Antwerp)
Fall 2017: Exchange program (Toulouse School of Management)

2018-22: PhD in Data Science (University of Antwerp | FWO funding)
Topic: Rule-based explanation methods to gain insight into classification models using big behavioral data
Spring 2022: Visiting researcher at Columbia Business School

2022-present: Data Scientist at QuantumBlack, AI by McKinsey

Icebreaker

What do you expect from this class?

What do you want to learn?

Law students' expectations according to ChatGPT

Understanding the Basics: Many law students may expect to gain a foundational understanding of AI, including key concepts, terminology, and how AI technologies work

Legal and Ethical Implications: They may expect discussions on topics such as AI bias and fairness, privacy concerns, liability issues, and the impact of AI on individual rights and societal values.

AI in Legal Practice: They may want to explore AI applications in tasks such as legal research, document analysis, contract review, and predictive analytics.

Regulatory Frameworks: Understanding the regulatory frameworks surrounding AI may also be important for law students.

Case Studies and Examples: Law students may appreciate case studies and real-world examples that illustrate how AI is being applied in various legal contexts.

Objectives for today

1. Introduction to Artificial Intelligence & Data Science
2. Deep dive: Predictive Modeling
3. Data Science for Business
4. Ethical Challenges

1. Introduction to AI & Data Science

“I believe that everyone should be communicating about Machine Learning, because it is the future. And there is too much glitter, and jargon, and science fiction in how these things are discussed today.”

Cassie Kozyrkov, (ex-)Chief Decision Scientist at Google

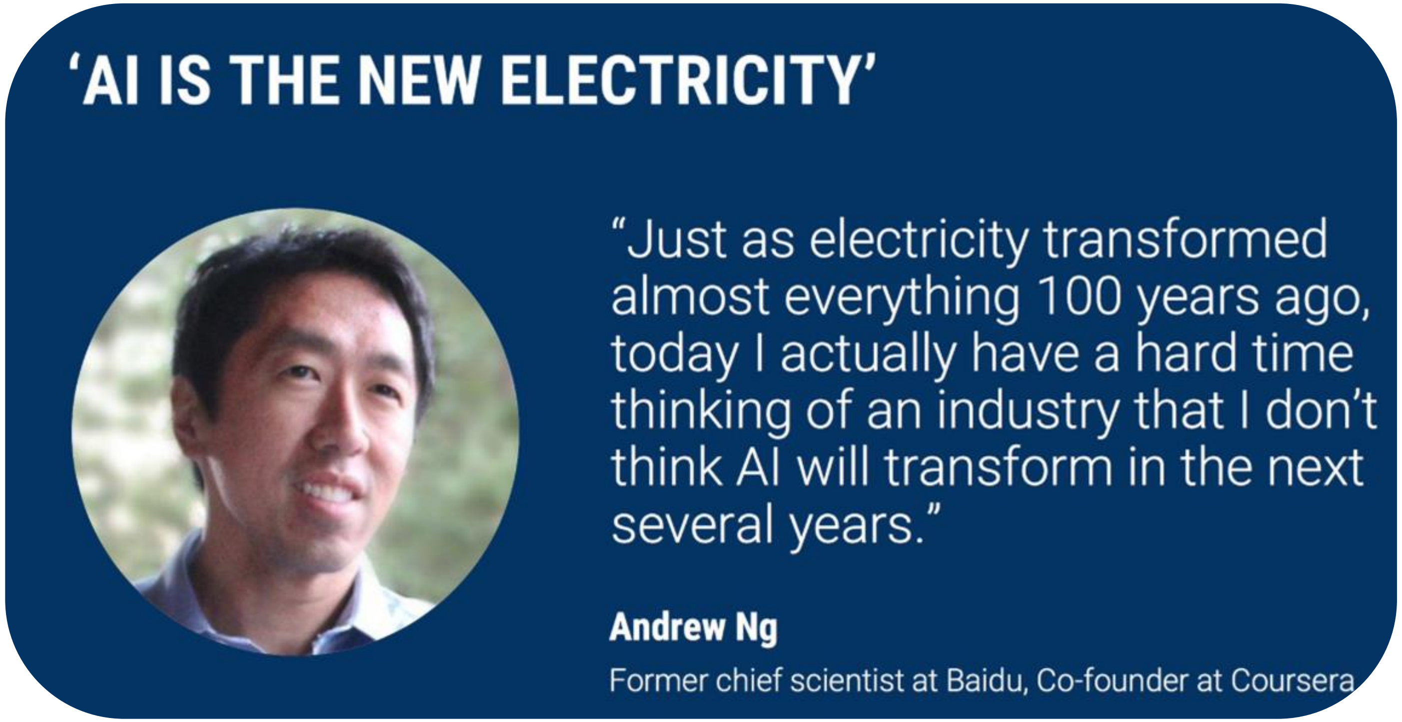

Terminology

Artificial Intelligence: make machines perform "intelligent" tasks on their own

Data Science: set of fundamental principles that guide extraction of knowledge from data

Data Mining: automatic extraction of patterns from large amounts of data (i.e., without human intervention or explicit programming)

Machine Learning: algorithms that incorporate Data Science principles to automatically learn patterns from large amounts of data

$y = f(x)$
If… then…

X1 = Age
X2 = Gender
X3 = Income
…

Data Mining

The **automatic** extraction of patterns from large amounts of data

Data Mining

The **automatic** extraction of patterns from large amounts of data

Motivation

Increase efficiency and/or accuracy (**benchmark** to size the value!)
More data and processing power lead to new opportunities

Motivation

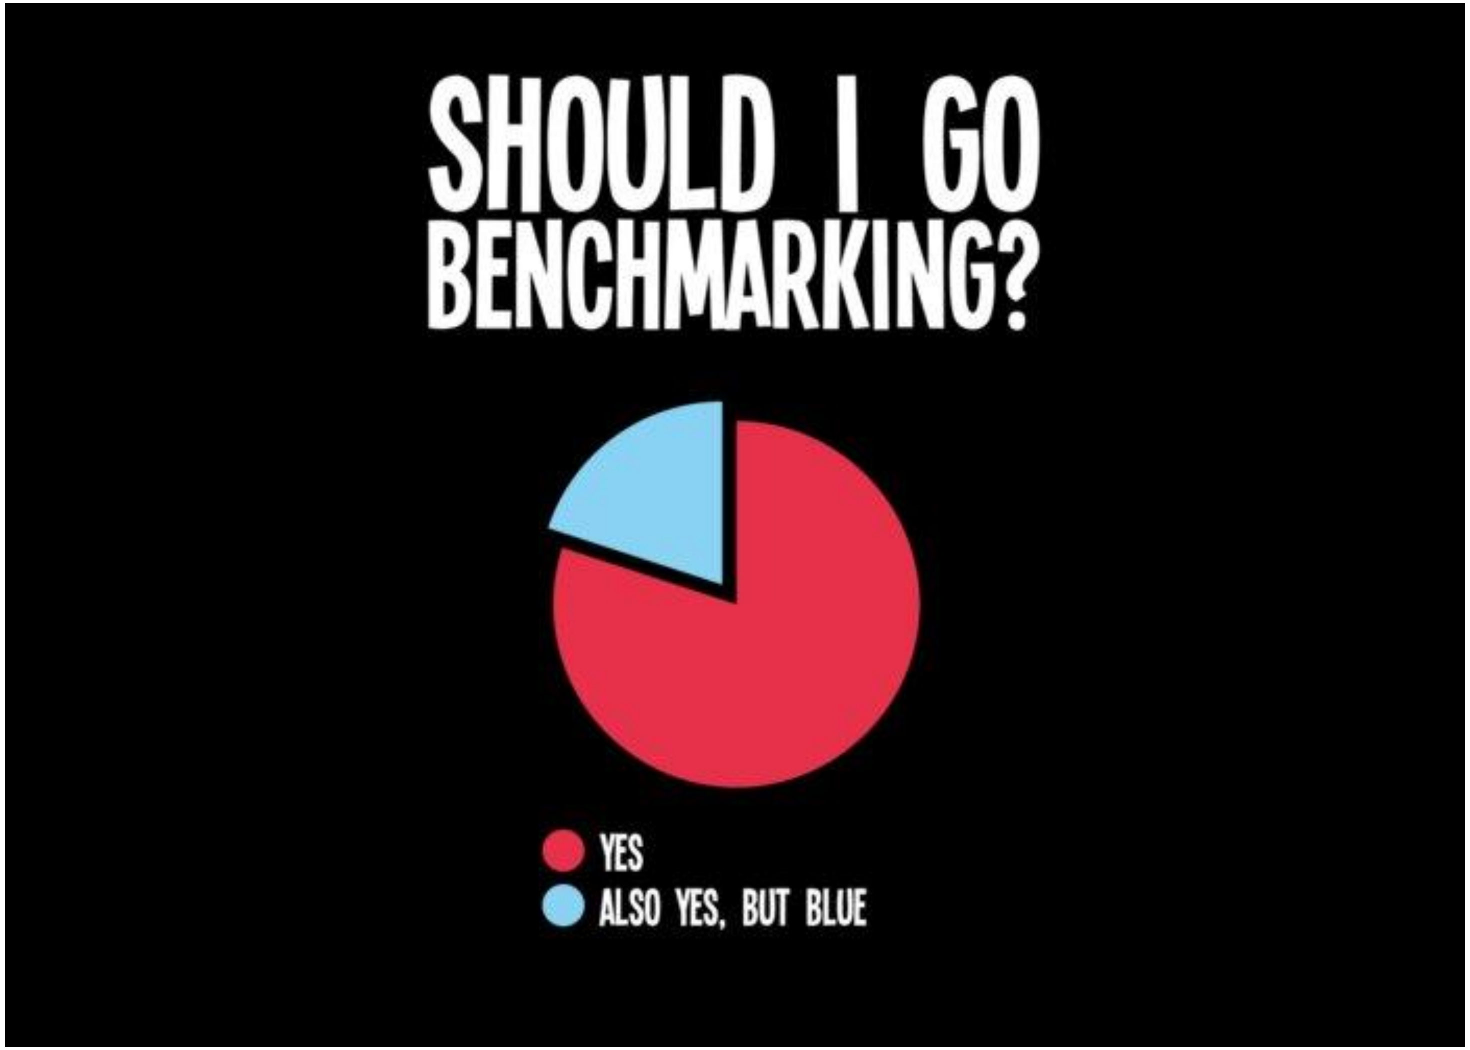

Augmenting human intelligence

AI to augment or replace human?

How to prioritize new AI usecases?

How to decide AI usecase is worth it?

…

- Bionic systems
- Human + AI
- Validate and test thoroughly
- Full automation for simple and low-risk tasks
- Do (incremental) benefits outweigh risks?

AI Should Augment Human Intelligence, Not Replace It

by David De Cremer and Garry Kasparov

[Harvard Business Review, 2021]

Not to be confused with

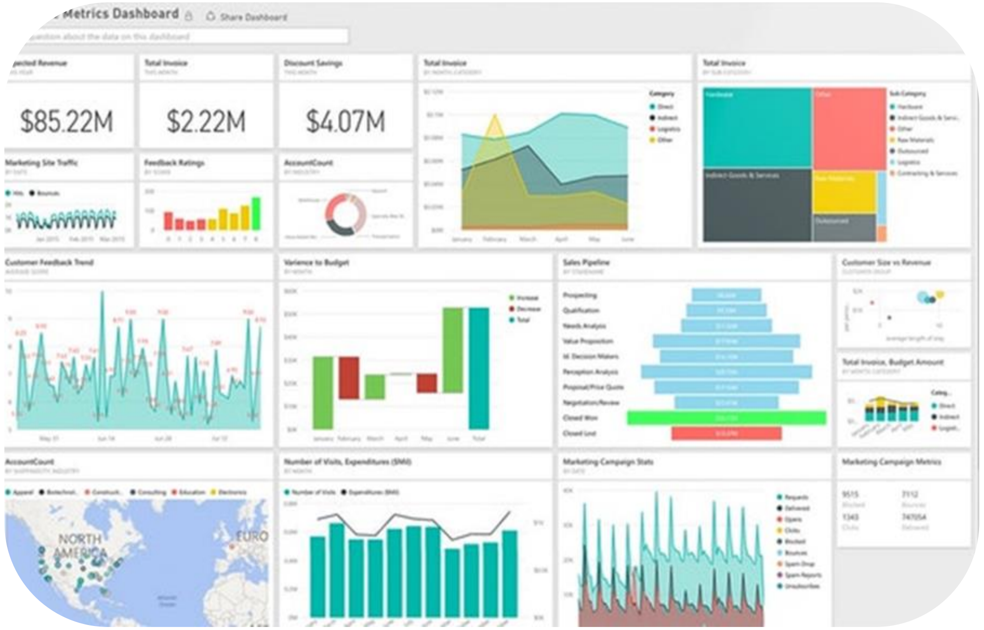

- **Querying/Reporting**

= users know exactly what they look for

- **Online Analytical Processing (OLAP)**

= advanced querying/reporting, multidimensional analysis, visualizations etc.

- **Business Intelligence (BI)**

= getting the right information to the right person at the right time

The end user is the engine of discovery!

Data

- **Tabular data:** low-dimensional, structured

"Features" or "variables"

Data "instances" or "examples"

	A	B	C	D	E	F	G	H	I	J
1	ID	Last Name	First Name	City	State	Gender	Student Status	Major	Country	Age
2	1	DOE01	JANE01	Los Angeles	California	Female	Graduate	Politics	US	30
3	2	DOE02	JANE02	Sedona	Arizona	Female	Undergraduate	Math	US	19
4	3	DOE01	JOE01	Elmira	New York	Male	Graduate	Math	US	26
5	4	DOE02	JOE02	Lackawana	New York	Male	Graduate	Econ	US	33
6	5	DOE03	JOE03	Defiance	Ohio	Male	Graduate	Econ	US	37
7	6	DOE04	JOE04	Tel Aviv	Israel	Male	Graduate	Econ	Israel	25
8	7	DOE05	JOE05	Cimax	North Carolina	Male	Graduate	Politics	US	39
9	8	DOE03	JANE03	Liberal	Kansas	Female	Undergraduate	Politics	US	21
10	9	DOE04	JANE04	Montreal	Canada	Female	Undergraduate	Math	Canada	18
11	10	DOE05	JANE05	New York	New York	Female	Graduate	Math	US	33
12	11	DOE06	JOE06	Hot Coffe	Mississippi	Male	Undergraduate	Econ	US	18
13	12	DOE06	JANE06	Java	Virginia	Female	Graduate	Math	US	38
14	13	DOE07	JOE07	Varna	Bulgaria	Male	Graduate	Politics	Bulgaria	30
15	14	DOE08	JOE08	Moscow	Russia	Male	Graduate	Politics	Russia	30
16	15	DOE07	JANE07	Drunkard Creek	New York	Female	Undergraduate	Math	US	21
17	16	DOE08	JANE08	Mexican Hat	Utah	Female	Undergraduate	Econ	US	18
18	17	DOE09	JANE09	Amsterdam	Holland	Female	Undergraduate	Math	Holland	19
19	18	DOE10	JANE10	Mexico	Mexico	Female	Graduate	Politics	Mexico	31
20	19	DOE11	JANE11	Caracas	Venezuela	Female	Undergraduate	Math	Venezuela	18
21	20	DOE09	JOE09	San Juan	Puerto Rico	Male	Graduate	Politics	US	33
22	21	DOE12	JANE12	Remote	Oregon	Female	Undergraduate	Econ	US	19
23	22	DOE10	JOE10	New York	New York	Male	Undergraduate	Econ	US	21
24	23	DOE13	JANE13	The X	Massachusetts	Female	Graduate	Politics	US	25
25	24	DOE14	JANE14	Beijing	China	Female	Undergraduate	Math	China	18
26	25	DOE11	JOE11	Stockholm	Sweden	Male	Undergraduate	Politics	Sweden	19
27	26	DOE12	JOE12	Embarrass	Minnesota	Male	Graduate	Econ	US	28
28	27	DOE13	JOE13	Intercourse	Pennsylvania	Male	Undergraduate	Math	US	20
29	28	DOE15	JANE15	Loco	Oklahoma	Female	Undergraduate	Econ	US	20
30	29	DOE14	JOE14	Buenos Aires	Argentina	Male	Graduate	Politics	Argentina	30
31	30	DOE15	JOE15	Acme	Louisiana	Male	Undergraduate	Econ	US	19

Data

- Tabular data: low-dimensional, structured
- **Big data:**
  - Volume, variety, velocity
  - Requires more storage and processing power
  - Data that is too large to be effectively processed by traditional methods or techniques

What happens in an Internet minute?

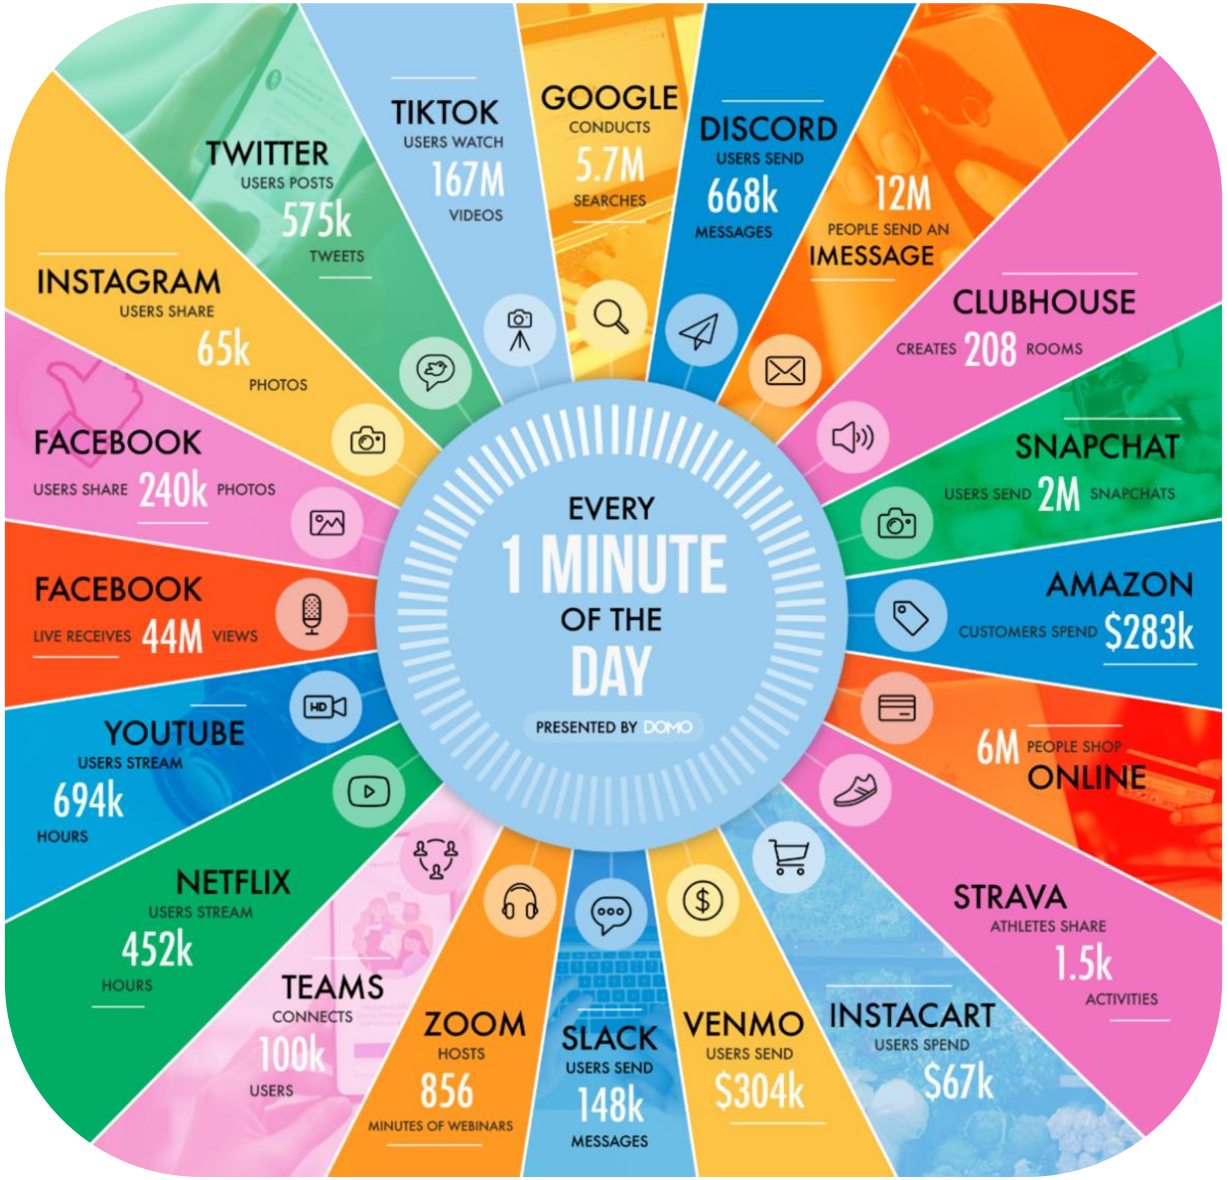

Source: Domo, 2021

Behavioral data

- Quality
- Zoom in
- Test causality (A/B testing)

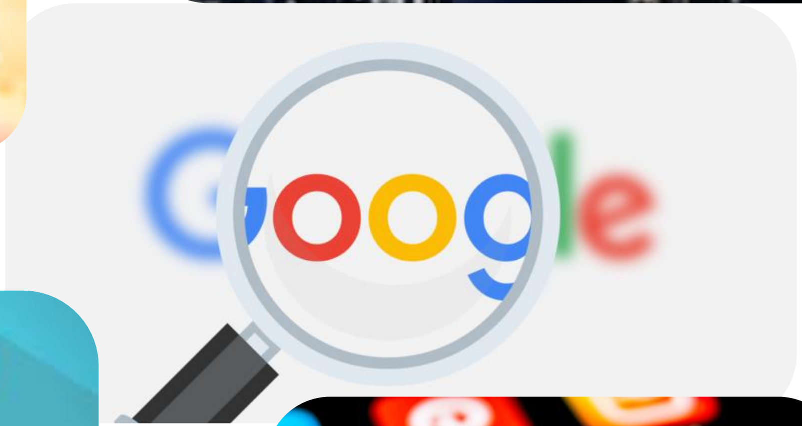

Source: [Seth Stephens-Davidowitz, Everybody Lies]

Behavioral data

- Quality
- Zoom in
- Test causality (A/B testing)

Cells	Cell 1 (Control)	Cell 2	Cell 3
Box Art	 Default artwork	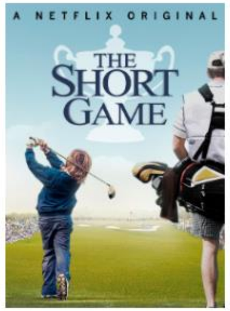 14% better take rate	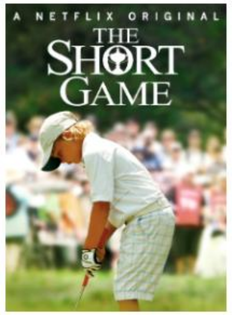 6% better take rate

[Netflix TechBlog, 3 May 2016]

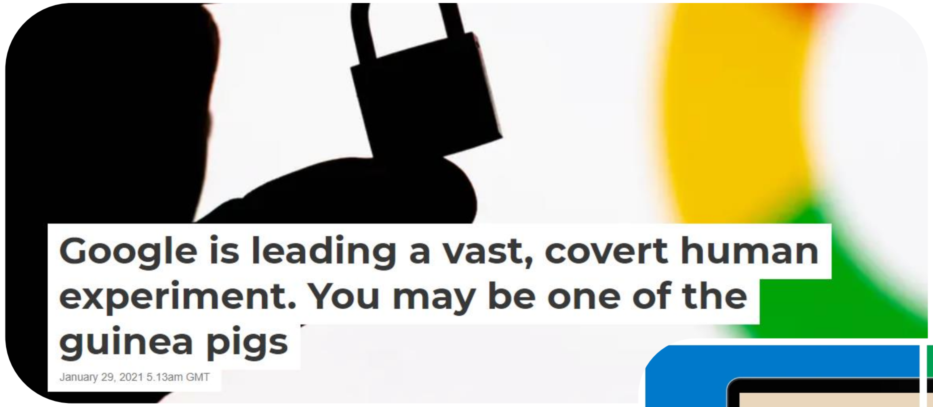

[The Conversation, 29 January 2021]

okcupid

[USA today news, 2014]

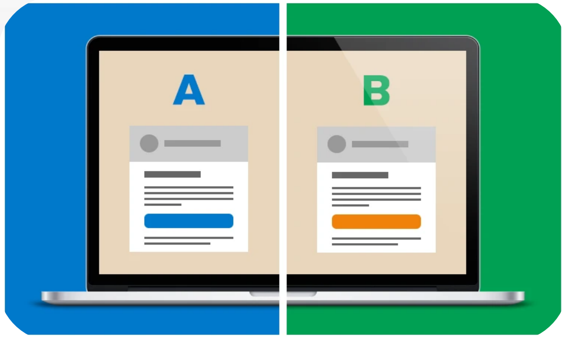

Source: [Seth Stephens-Davidowitz, Everybody Lies]

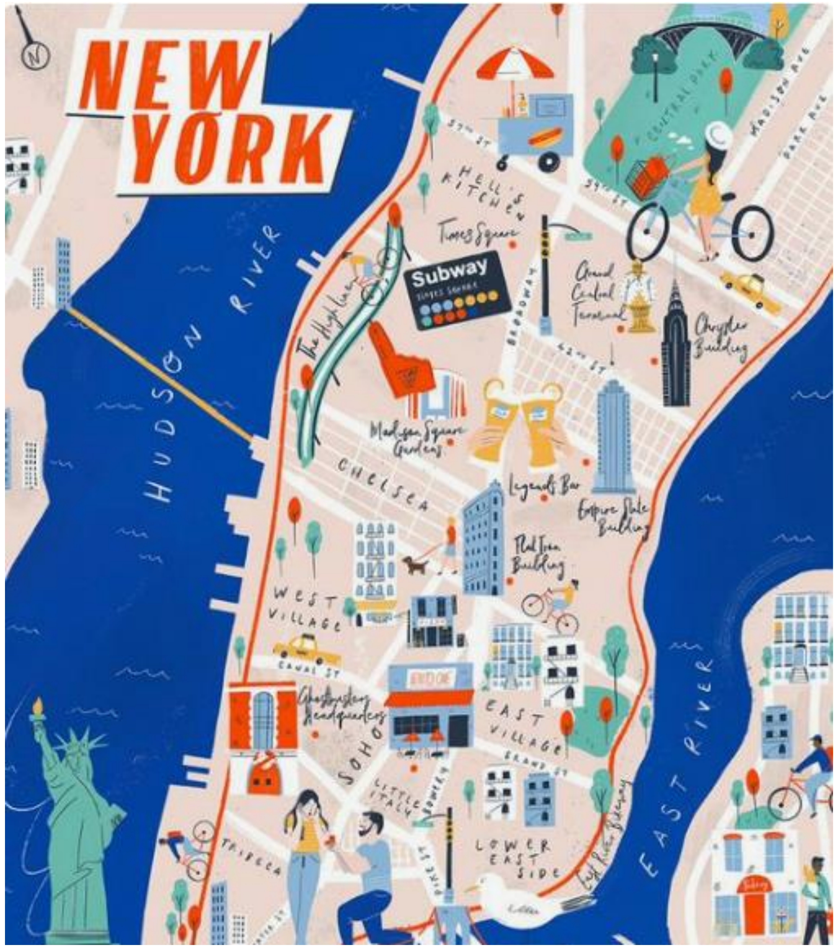

Dimension of the evidence pool: 50,000

Persons

Name	Time Square 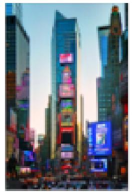	Dumbo 	...	Columbia University 	# evidence present (last month)	*Model decision: Tourist?*
...	...	...	...	...	...	...
Anna	1	1	...	0	85	*YES*
Jack	0	0	...	1	129	*NO*
Billie	1	1	...	1	53	*YES*
...	...	...	...	...	...	...

Figure 1: Fictitious example of location data used to discriminate tourists and NY citizens

Source: [Ramon Y, Evidence counterfactuals, KDNuggets blogpost, 2020]

Terminology

Artificial Intelligence: make machines perform "intelligent" tasks on their own

Data Science: set of fundamental principles that guide extraction of knowledge from data

Data Mining: automatic extraction of patterns from large amounts of data (i.e., without human intervention or explicit programming)

Machine Learning: algorithms that incorporate Data Science principles to learn automatically learn patterns from large amounts of data

Deep Learning: subset of ML algorithms that use a brain-like logical structure called "Artificial Neural Networks" to learn patterns from large amounts of (unstructured) data

Terminology

Deepdive on next slide

Artificial Intelligence: make machines perform "intelligent" tasks on their own

Data Science: set of fundamental principles that guide extraction of knowledge from data

Data Mining: automatic extraction of patterns from large amounts of data (i.e., without human intervention or explicit programming)

Machine Learning: algorithms that incorporate Data Science principles to learn automatically learn patterns from large amounts of data

Deep Learning: subset of ML algorithms that use a brain-like logical structure called "Artificial Neural Networks" to learn patterns from large amounts of (unstructured) data

We perceive the world through our senses

Machines perceive the world through data

Structured enterprise data

10-20%

High degree of organization, easy to search

Unstructured enterprise data

80-90%

Information that is difficult to organize or search

Examples of unstructured data

Image

Video

Audio

Text

Webpage

More data and processing power unlocks the potential of Deep Learning in business and society

Exercise: Describe the use case and impact of the Deep Learning applications below. What data source is used as input for each application?

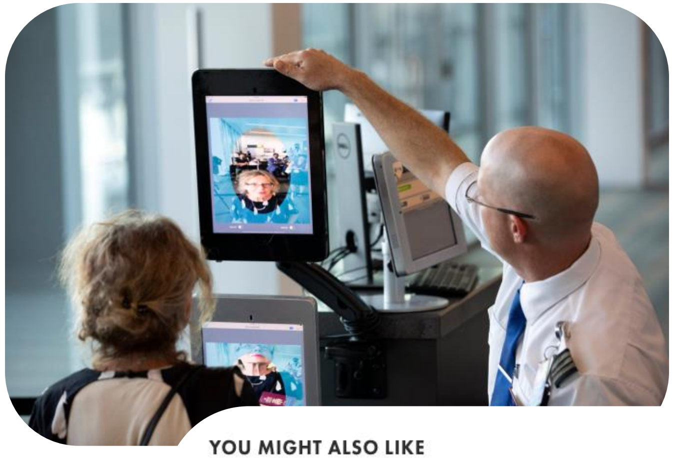

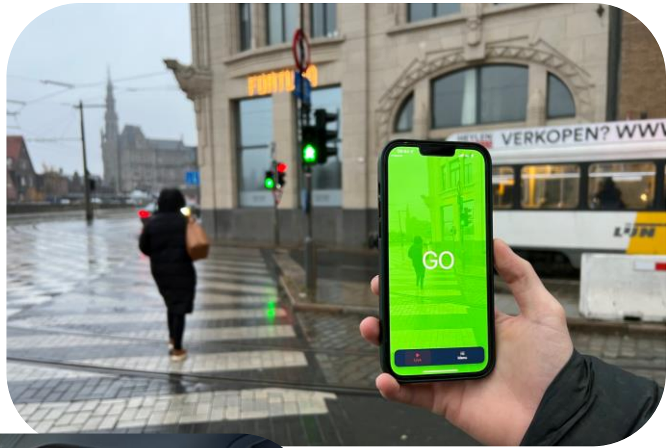

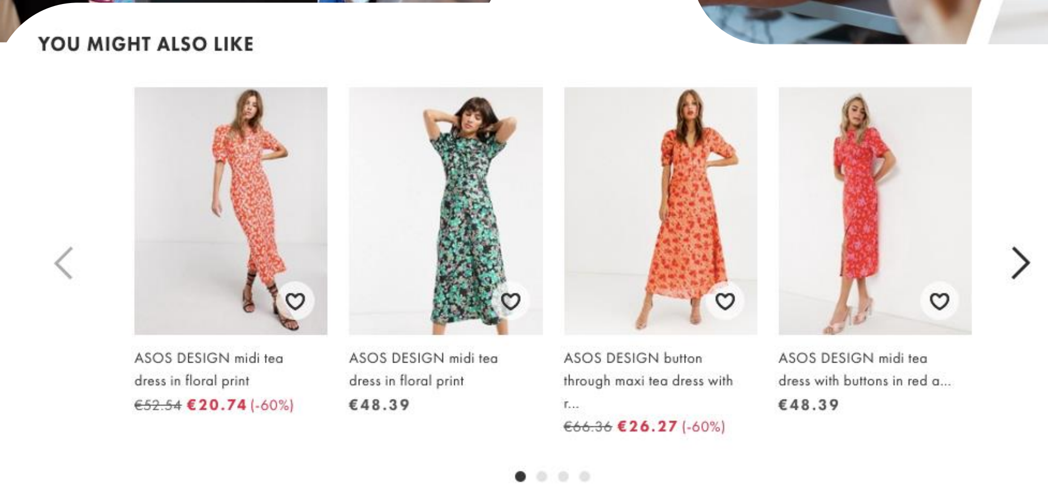

Gen AI vs. Discriminative AI

- **Generative AI (Gen AI)** = subset of Deep Learning
- Gen AI generates new data based on a known distribution of existing data (e.g., text) to learn what the expected outcome should be. Uses prompts as inputs.
  - Examples: ChatGPT (text to text), DreamStudio (text to image), ...
  - Output: natural language, image, video, audio
- **Discriminative AI** = classifies or predicts outcome of interest for new data based on historical (labelled) data set
  - Examples: Image classification, targeted advertising model, demand forecasting, ...
  - Output: class, discrete number, score, probability

Examples of Data Mining tasks

Frequent Itemsets

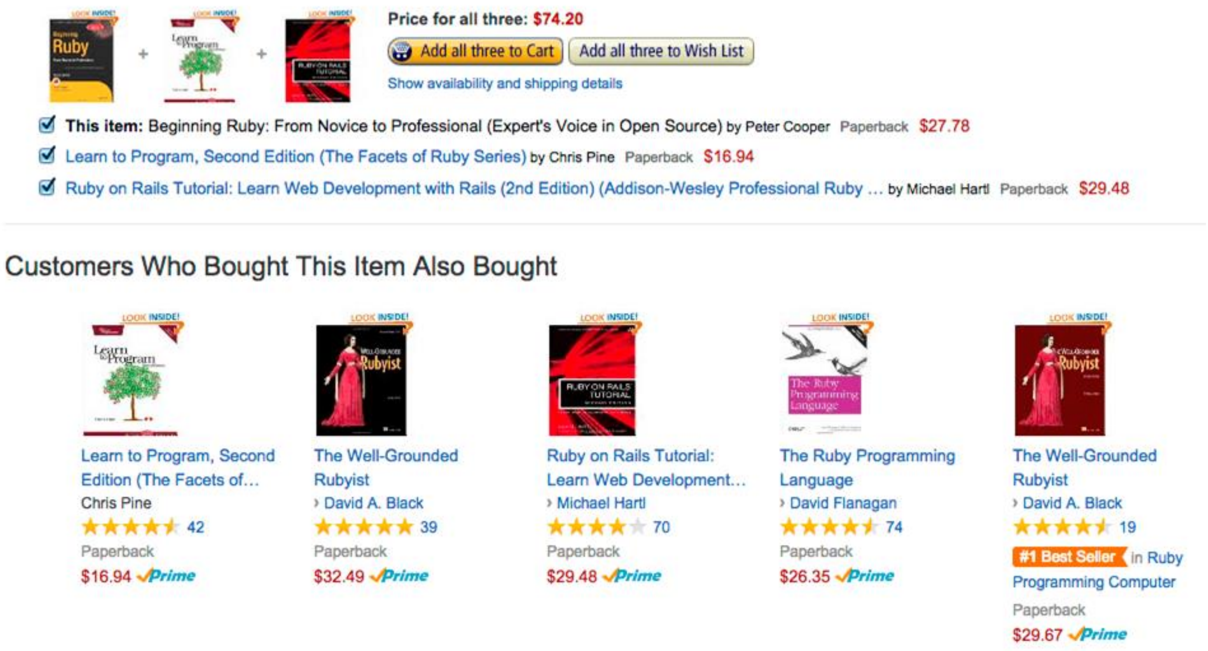

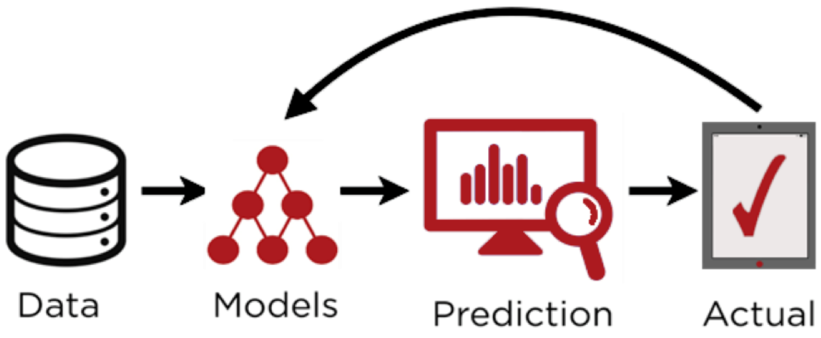

Predictive Modeling (i.e., Discriminative AI)

Recommender Systems

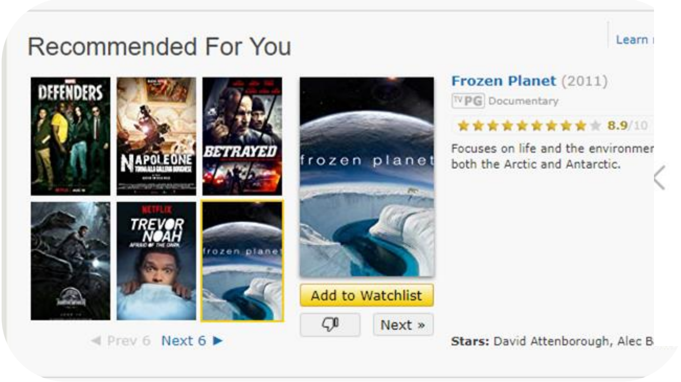

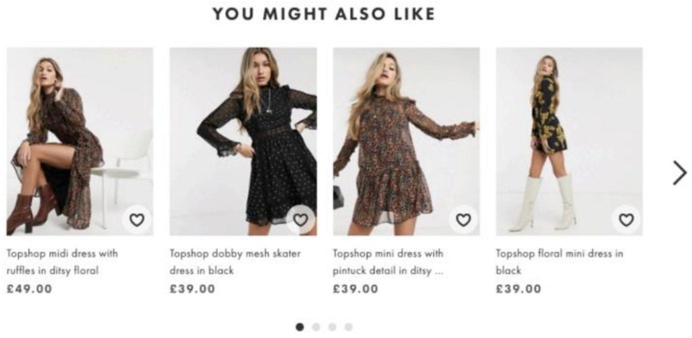

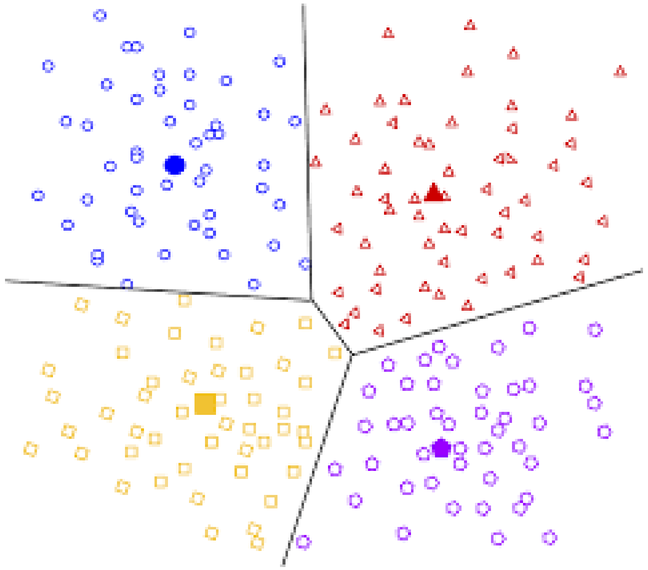

Clustering

2. Predictive modeling

Input data (X)	Target variable (Y)
Socio-demographic data	Product interest
Social media (e.g., Facebook likes)	Political preference
Criminal record	Recidivism
Images	Categories
Web page visits	Product interest

Most of AI is about extracting or “learning” this mapping (pattern) from (historical) data

= “Supervised learning”





Example: credit scoring

- Banks: should I grant credit to this loan applicant?
- Predict the creditworthiness based on historical data

Data Mining process

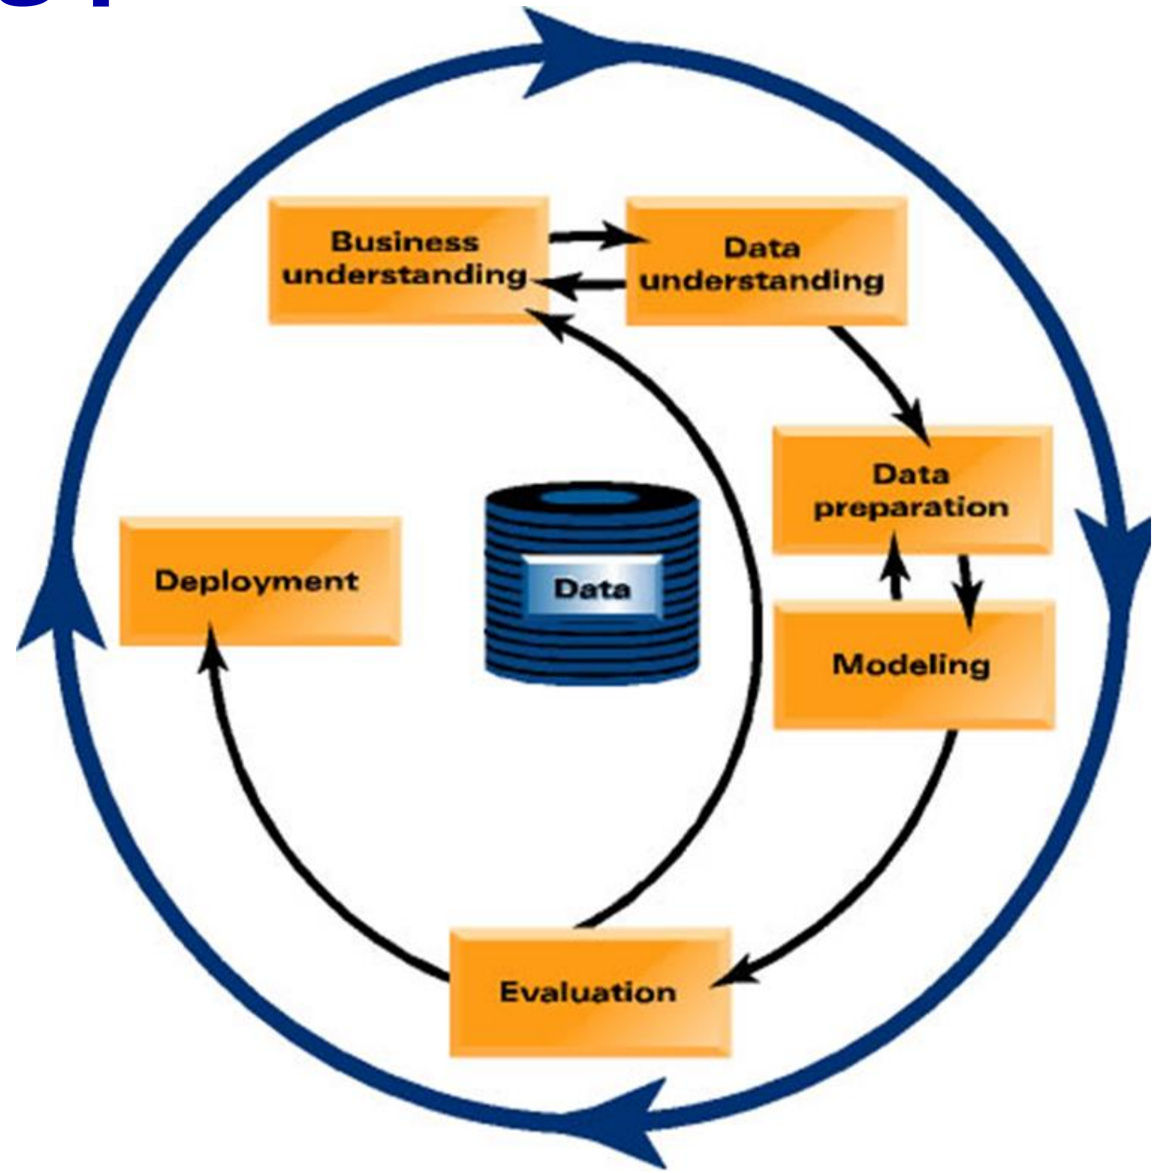

Source: Cross-industry Standard Process for Data Mining (CRISP-DM)

Example: credit scoring

- Banks: should I grant credit to this loan applicant?
- Predict the creditworthiness based on historical data

Client	Income	Sex	Amount	Default
A	1,600	M	175,000	N
B	2,000	F	350,000	Y
C	3,500	M	55,000	N
D	950	M	129,000	Y
E	10,500	M	1,000,000	N
F	5,700	F	240,000	N
G	2,400	F	250,000	N

Data

Data mining technique

Pattern

Prediction model:

If Income<2,500 and Amount>100,000, then: Default = Yes
Else: Default = No

Client	Income	Sex	Amount	Default
New client	2,000	F	500,000	**?**

Note: an initial set of data instances with **known** target variable needed!

Data preprocessing

- Dirty, noisy data (e.g., income = -156)
- Inconsistent data (e.g., purpose = car, mortgage = yes)
- Data integration issues (e.g., amount in Euro and USD)
- Missing value treatment (e.g., income = ?)
- Duplicate data (e.g., salary vs. income)
- Outlier treatment
- Normalising
- Etc.

→ Time-consuming (80% rule)

→ Garbage in = Garbage out

Pattern

- **Mathematical models (linear vs. nonlinear)**
- Rule-based models

$f(x) > 0.5$ => customer = good
$f(x) \leq 0.5$ => customer = bad

- Linear
  - Linear, logistic regression; linear discriminant analysis
  - Result: linear function of attributes
  - $f(x) = 0.125$ income $+ 0.305$ age $- 0.02$ gender $+ \ldots + 3.1$ amount loan
- Non-linear
  - Artificial Neural Networks, Support Vector Machines, RVM, ...
  - Result: non-linear function of attributes
  - $f(x) = 0.201$ income^2 age^3 $- 0.55$ age^3 $- 5.21$ gender income $+ \ldots + 3.6$ gender^2 amount loan^2

Pattern

- Mathematical models (linear vs. nonlinear)
- **Rule-based models**

```
if (Checking Account < 200 DM and Duration > 15 m and
    Credit History = no credits taken and Savings Account < 1000 DM)
then class = bad
else if (Purpose = new car/repairs/education/others and
    Credit History = no credits taken/all credits paid back duly at this bank and
    Savings Account < 1000 DM)
then class = bad
else if (Checking Account < 0 DM and
    Purpose = furniture/domestic appliances/business and
    Credit History = no credits taken/all credits paid back duly at this bank and
    Savings Account < 500 DM)
then class = bad
else if (Checking Account < 0 DM and Duration > 15 m and
    Credit History = delay in paying off in the past and
    Savings Account < 500 DM)
then class = bad
else class = good
```

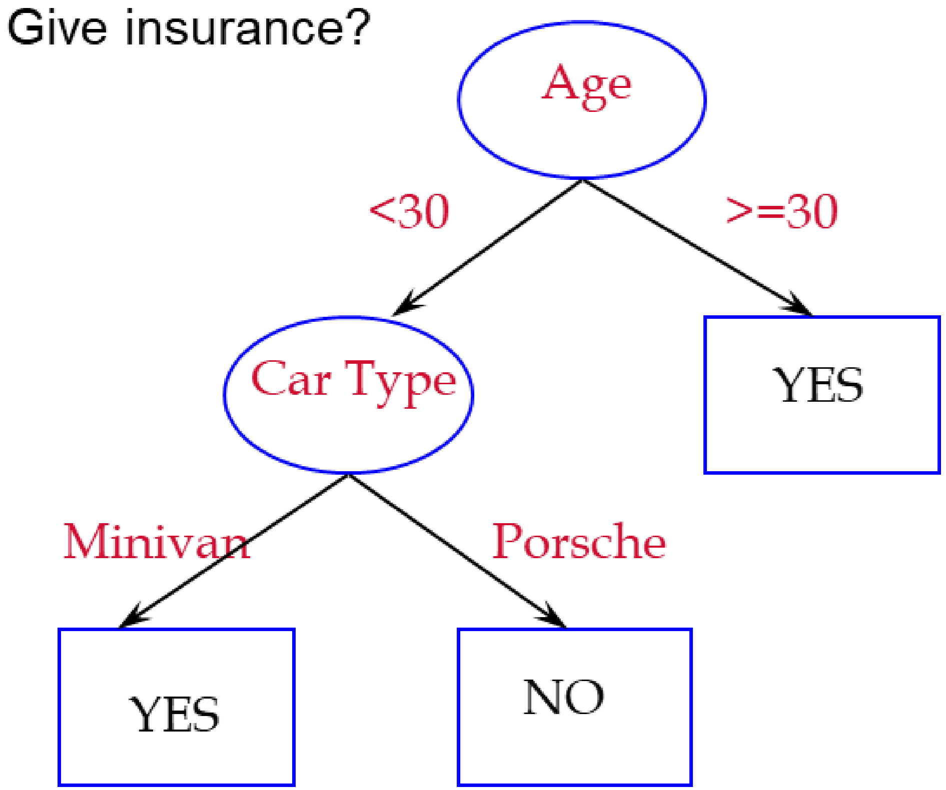

Pattern

- Mathematical models (linear vs. nonlinear)
- Rule-based models

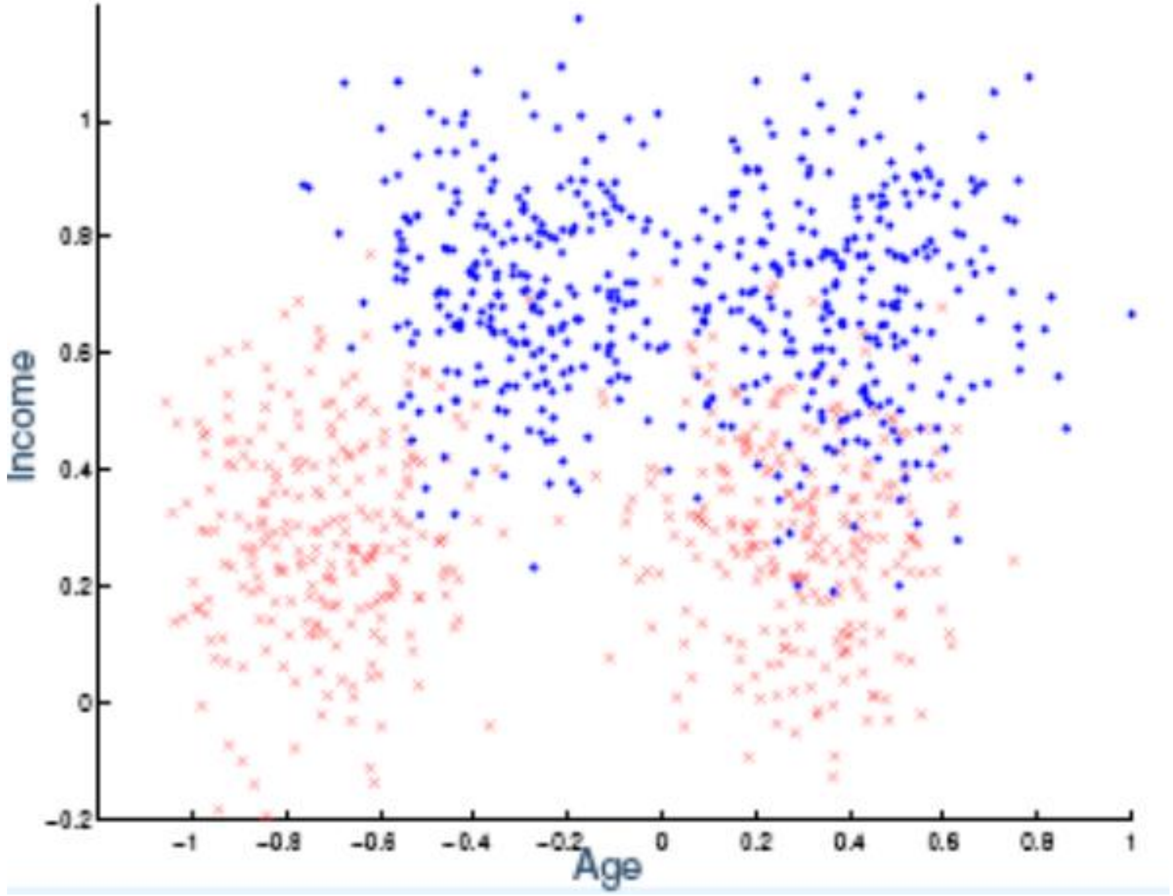

Pattern

- Mathematical models (linear vs. nonlinear)
- Rule-based models

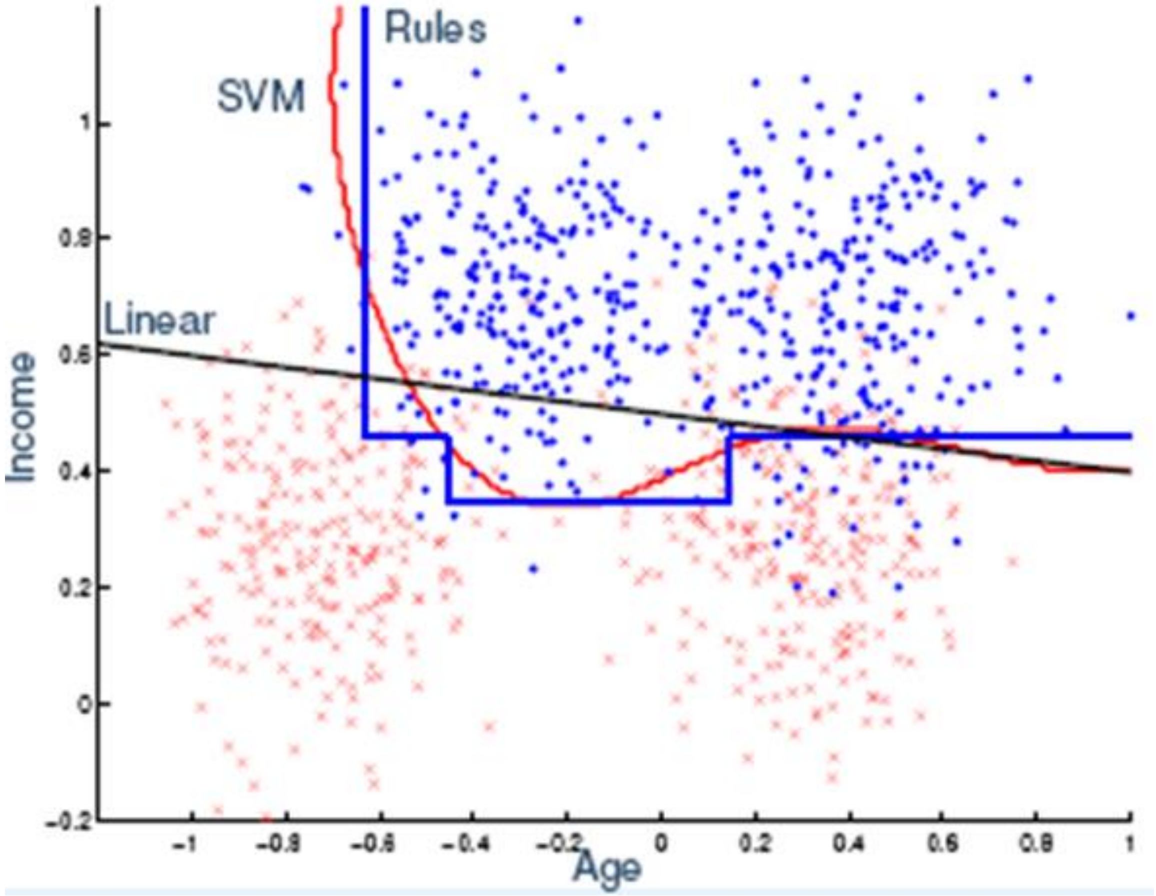

Data Mining techniques

- Different techniques for different patterns:
  - Linear: Linear regression, LDA, Logit,...
  - Non-linear: Deep Learning, Support vector machines, Random Forest,...
  - Rule-based: C4.5, RIPPER, CART,...
- Apply concept of finding informative variables

"So, if an old pirate gives me information about where a treasure is hidden..."

Data Mining techniques

Finding informative variables: Intuitively, which split is best?

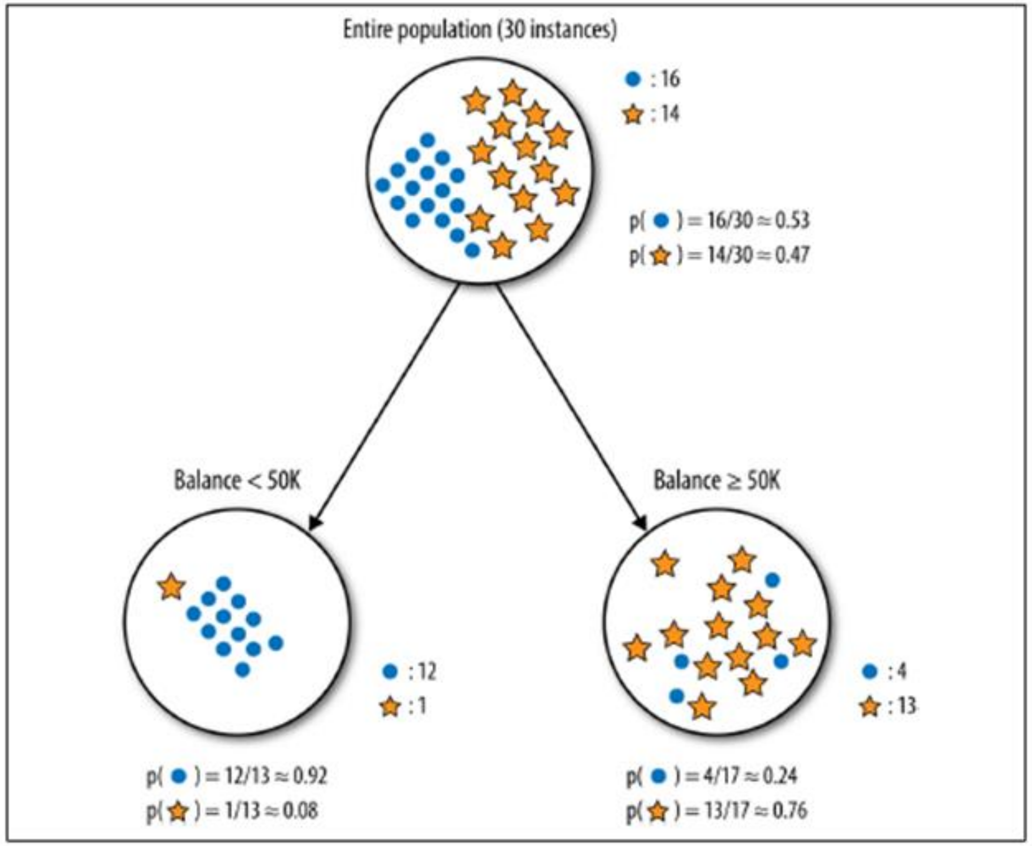

Figure 3-4. Splitting the "write-off" sample into two segments, based on splitting the Balance attribute (account balance) at 50K.

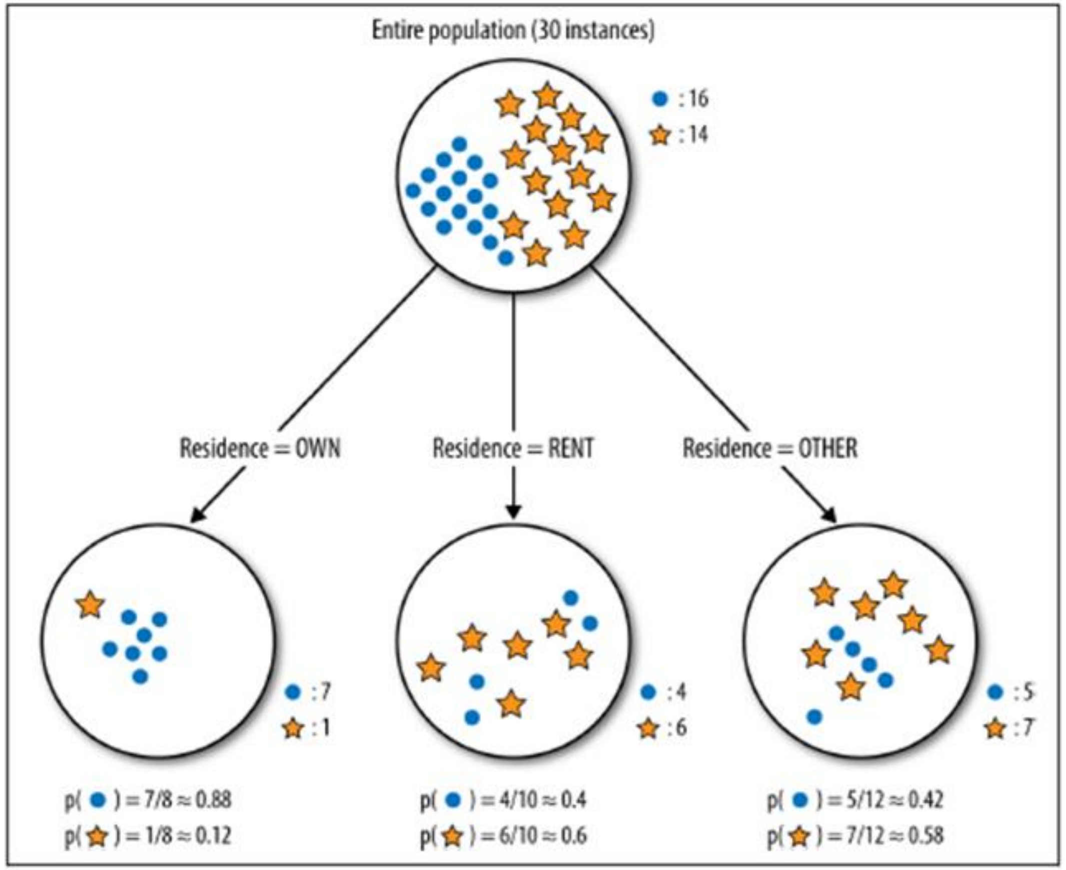

Figure 3-5. A classification tree split on the three-valued Residence attribute.

Evaluation

- Goal is to obtain models that **generalize well** to new observations
- Different metrics (e.g., accuracy = % of correct predictions)
- Create lab test of generalization performance (“test” vs. “training” data)

Deployment

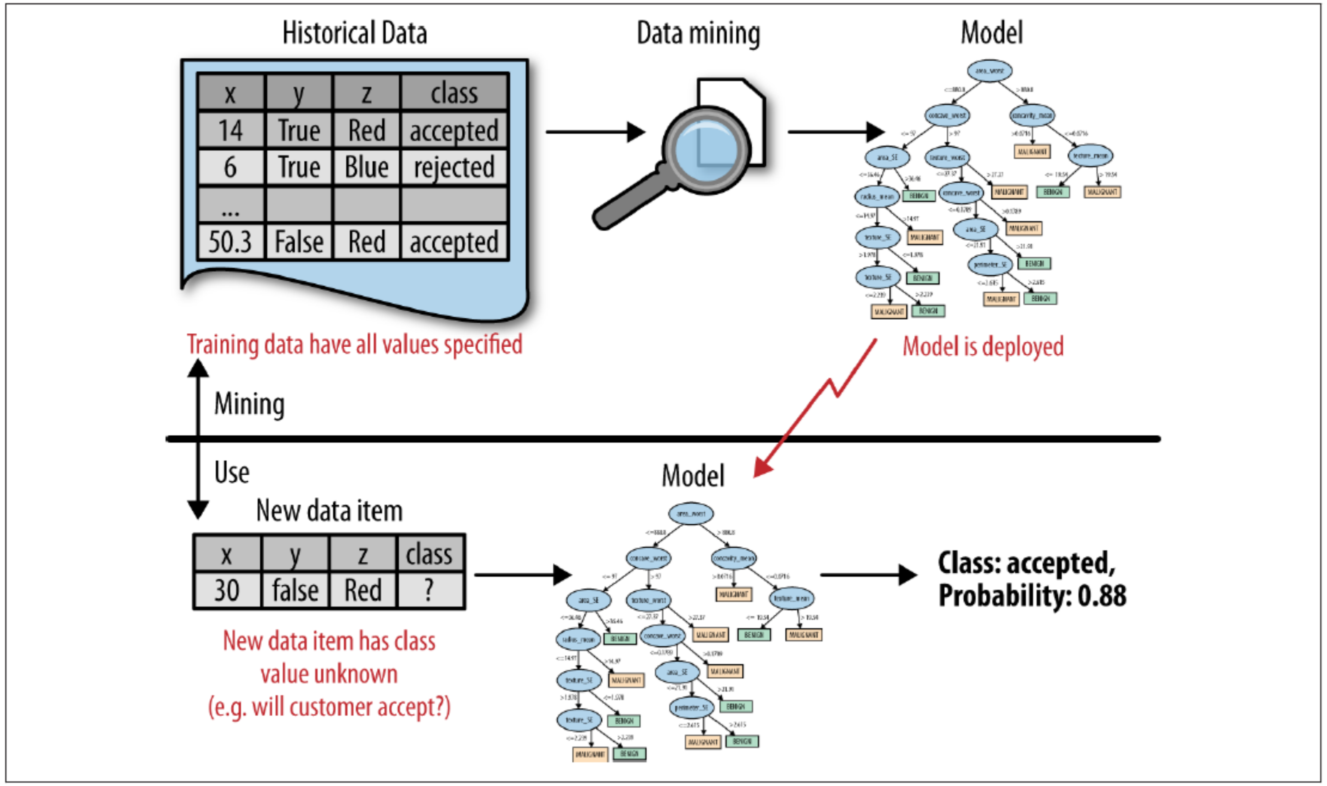

Figure 2-1. Data mining versus the use of data mining results. The upper half of the figure illustrates the mining of historical data to produce a model. Importantly, the historical data have the target ("class") value specified. The bottom half shows the result of the data mining in use, where the model is applied to new data for which we do not know the class value. The model predicts both the class value and the probability that the class variable will take on that value.

Note: an initial set of data instances with **known** target variable needed!

3. Data Science for Business

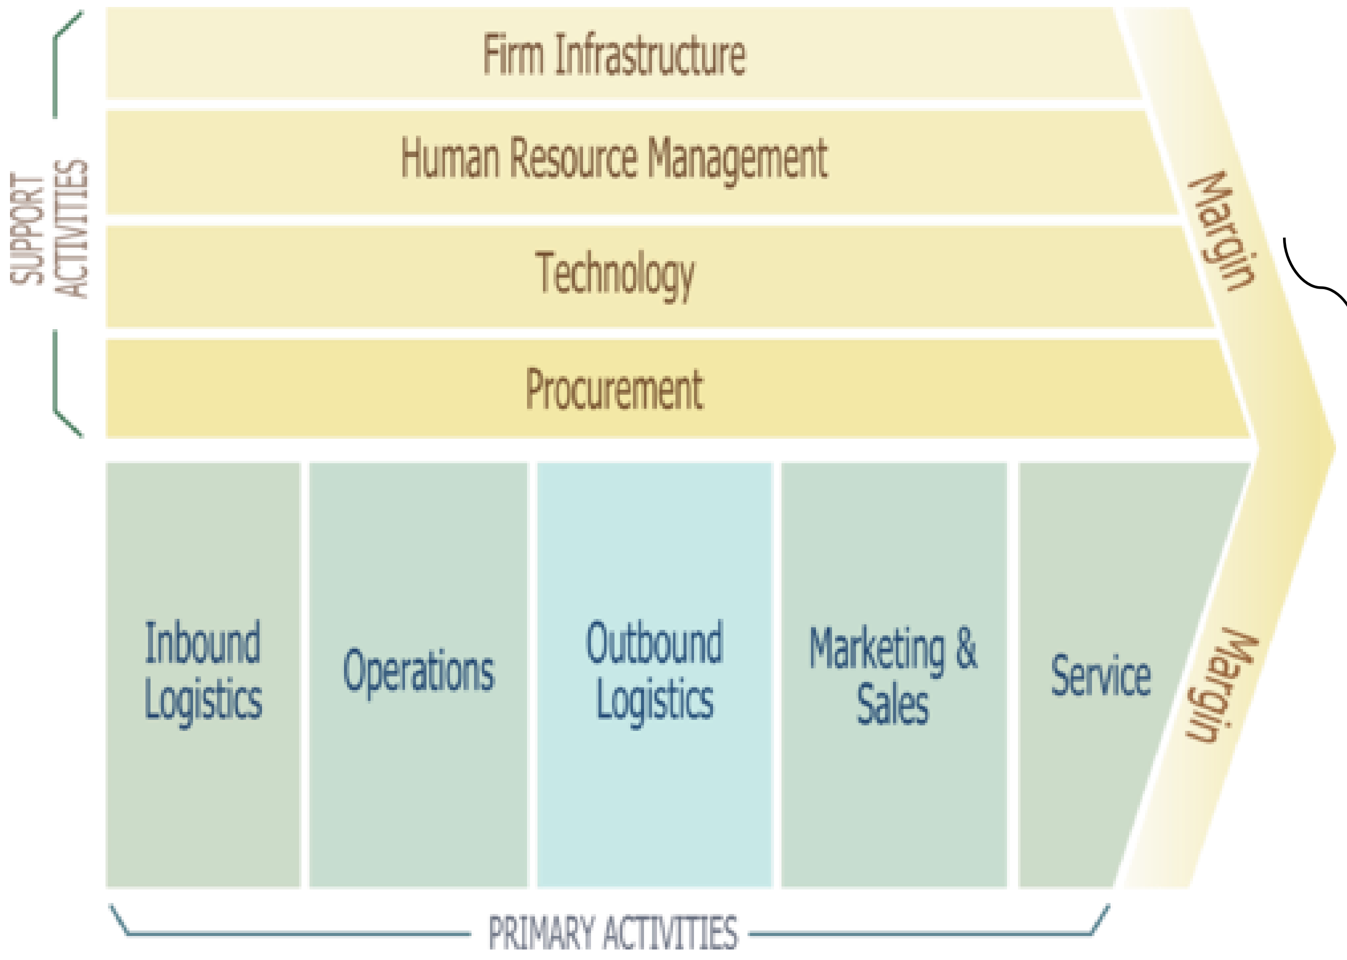

Profit = Revenue - Cost

Source: [Wikipedia, Michael Porter's value chain]()

Targeted advertising

- Optimize targeting, budgeting and audiences to drive growth
- Mini-case: Predicube (now: Pebblemedia)
- Data Mining task: Predictive modeling

Targeted advertising

- **Input data:**
  - Billions of records of persons visiting web pages
  - Tens of millions of unique cookie IDs
  - 20 million unique URLs
  - Daily number crunching on 50 servers for 5 hours (Amazon Web Services)
- **Data Mining task:** Find patterns in web browsing behavior

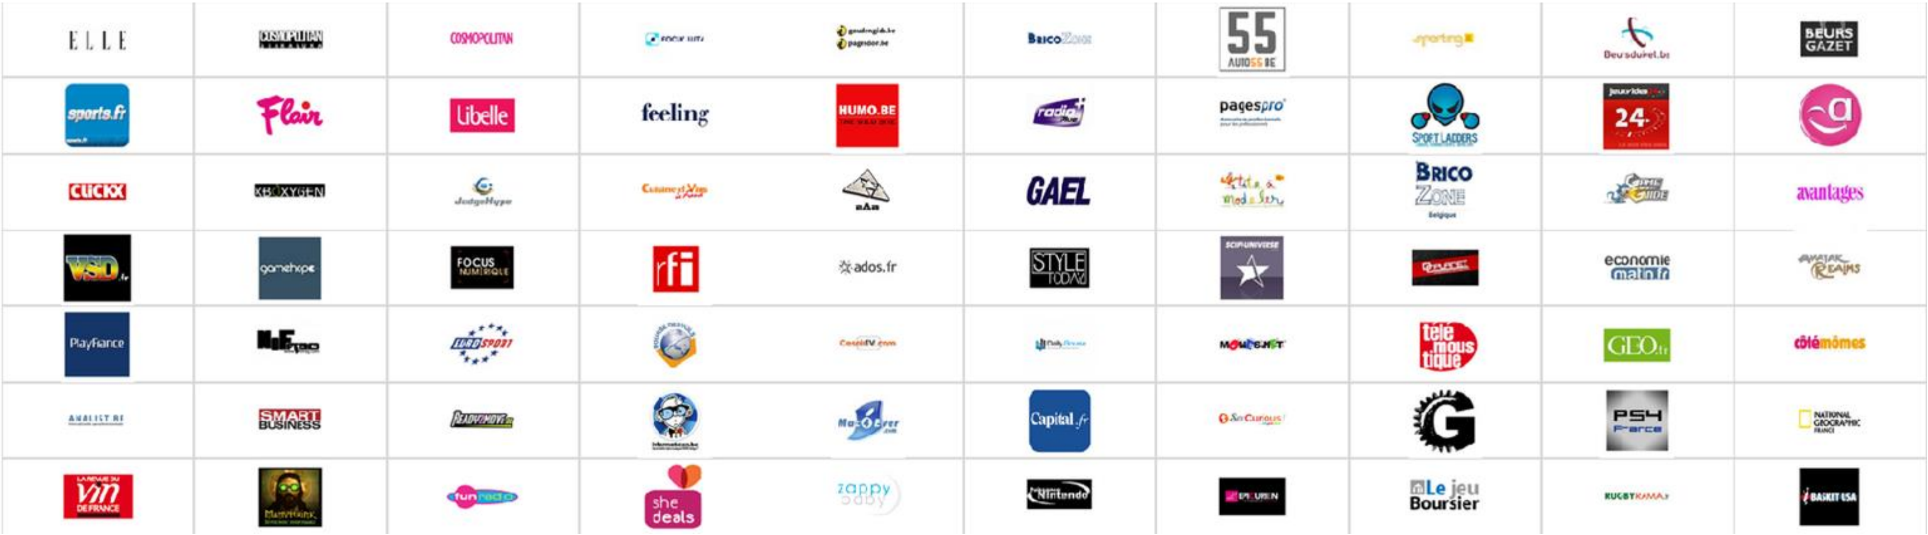

Find patterns in web browsing behavior

1. Once we know the browsing and conversion behavior of the existing customers (those who visited the homepage or conversion page of the advertiser)
2. We can identify specific pattern in their historical web click (and other) data
3. Find similar segments of prospects who show the same behavior
4. Show targeted ad to these “act-a-likes” of the known customers

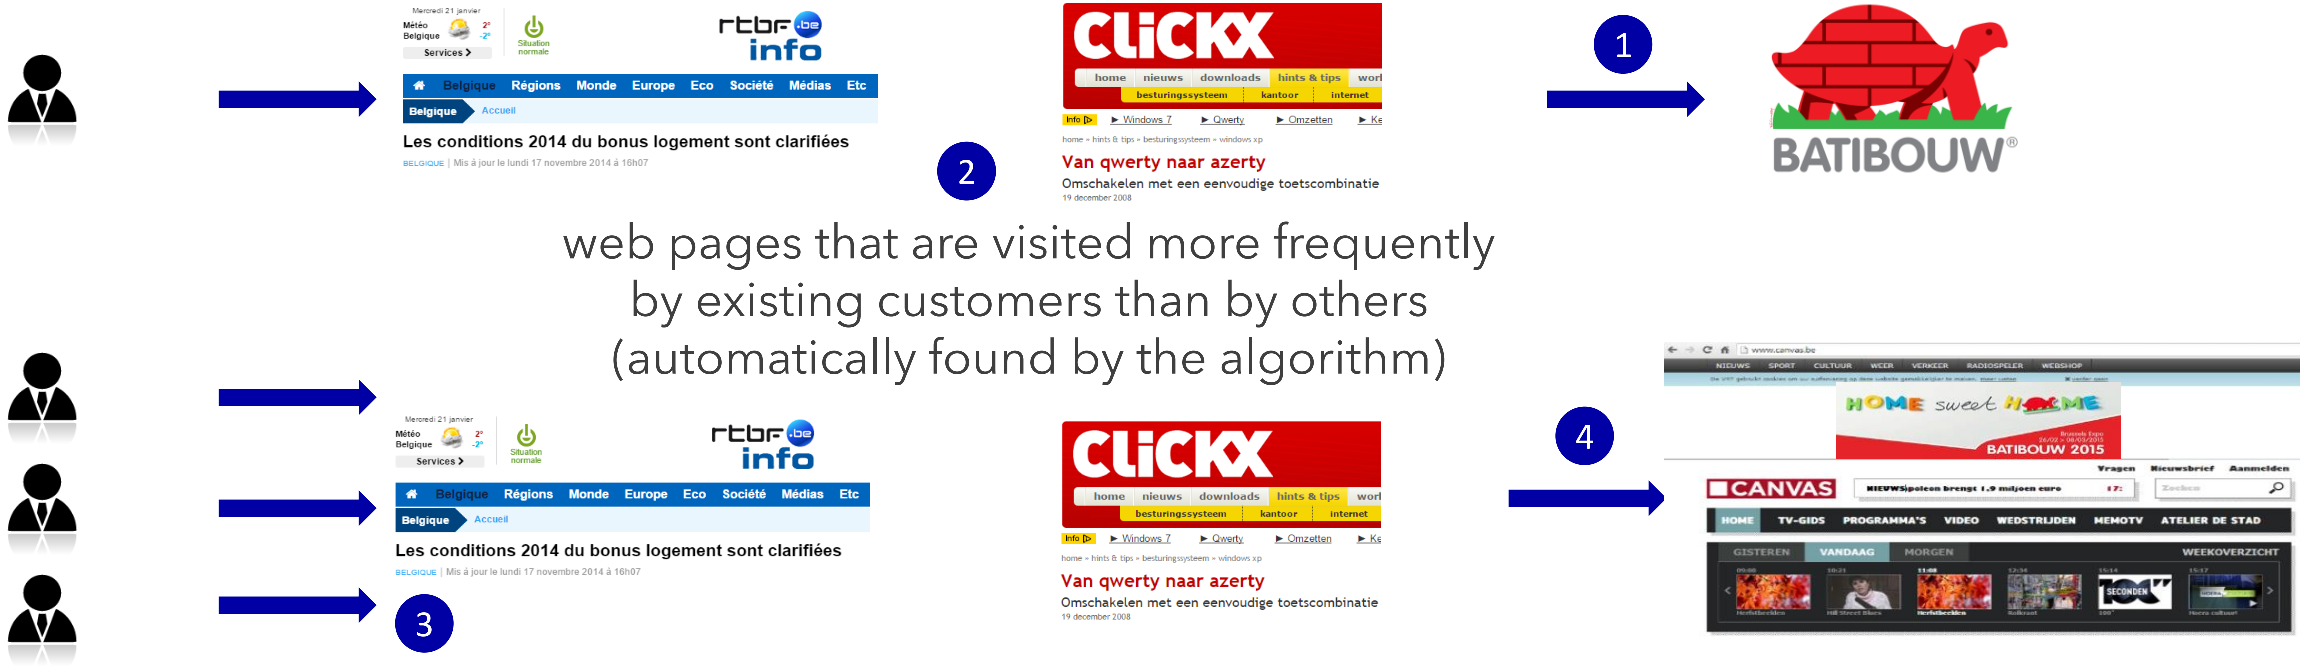

Find patterns in web browsing behavior

What does one data instance represent?
What is the target variable **Y**?
What are the input variables **xj**? What are potential values for the input variables?
What is the set of data instances with known values for the target variable?
How would you evaluate this data-driven targeting system?

Customer segmentation

- Know Your Customer (KYC) informs pricing strategy, distribution channels, communication, advertising campaigns, product & market development,...
- Mini-case: Oatly
- Data Mining task: Clustering

Customer segmentation

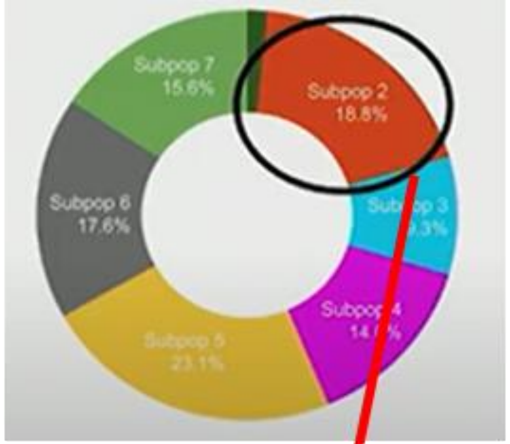

TOP FEATURES

Domains	Category
www.veganyackattack.com	Vegan recipe blog
www.therealllife-rd.com	Nutrition blog
www.sweetsimplevegan.com	Vegan recipe blog
www.fullofplants.com	Vegan recipe blog
www.elephantasticvegan.com	Vegan recipe blog
www.yourveganmom.com	Vegan recipe blog
www.kaleandcaramel.com	Recipe blog
www.vegannie.com	Vegan recipe blog
www.veggiesdontbite.com	Vegan recipe blog
www.thecolorfulkitchen.com	Vegan recipe blog
www.greenevi.com	Vegan recipe blog
www.plantpoweredkitchen.com	Vegan recipe blog
www.thevietvegan.com	Vegan recipe blog
www.katalysthealthblog.com	Vegan recipe blog
www.withfoodandlove.com	Recipe blog
www.keepinitkind.com	Vegan recipe blog
www.tworaspberries.com	Vegan recipe blog
www.tworaspberries.com	Vegan recipe blog

Source: [Presentation Melinda Williams (Dstillery), March 2019, STRATA Conference]

Personalization

- Products/services, user interface, pricing,...
- Mini-case: ASOS
- Data Mining task: Predictive modeling

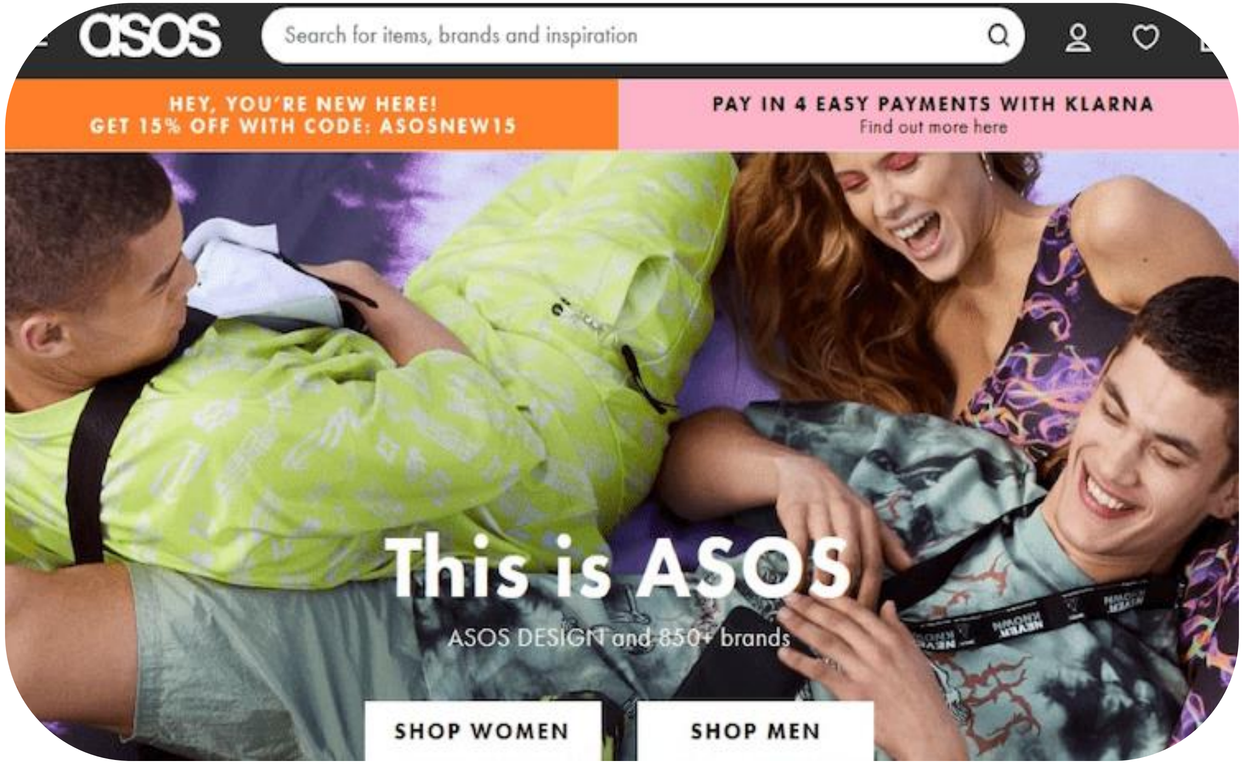

Personalization

- Recommendation problem:
  - ❑ Product catalogue: information overload (>85,000 products)
  - ❑ How can they help customers find products in overwhelming catalogue?
- Solutions: Home page recommendations, YMAL (product page recommendations), brand recommendations, outfit recommendations

Personalization

- Recommendation problem:
  - ❑ Product catalogue: information overload (>85,000 products)
  - ❑ How can they help customers find products in overwhelming catalogue?
- Solutions: Home page recommendations, YMAL (product page recommendations), brand recommendations, **outfit recommendations**

"What shoes will go with this dress?"

"What can I wear to a party?"

"Which items should I add to my wardrobe for summer?"

Personalization

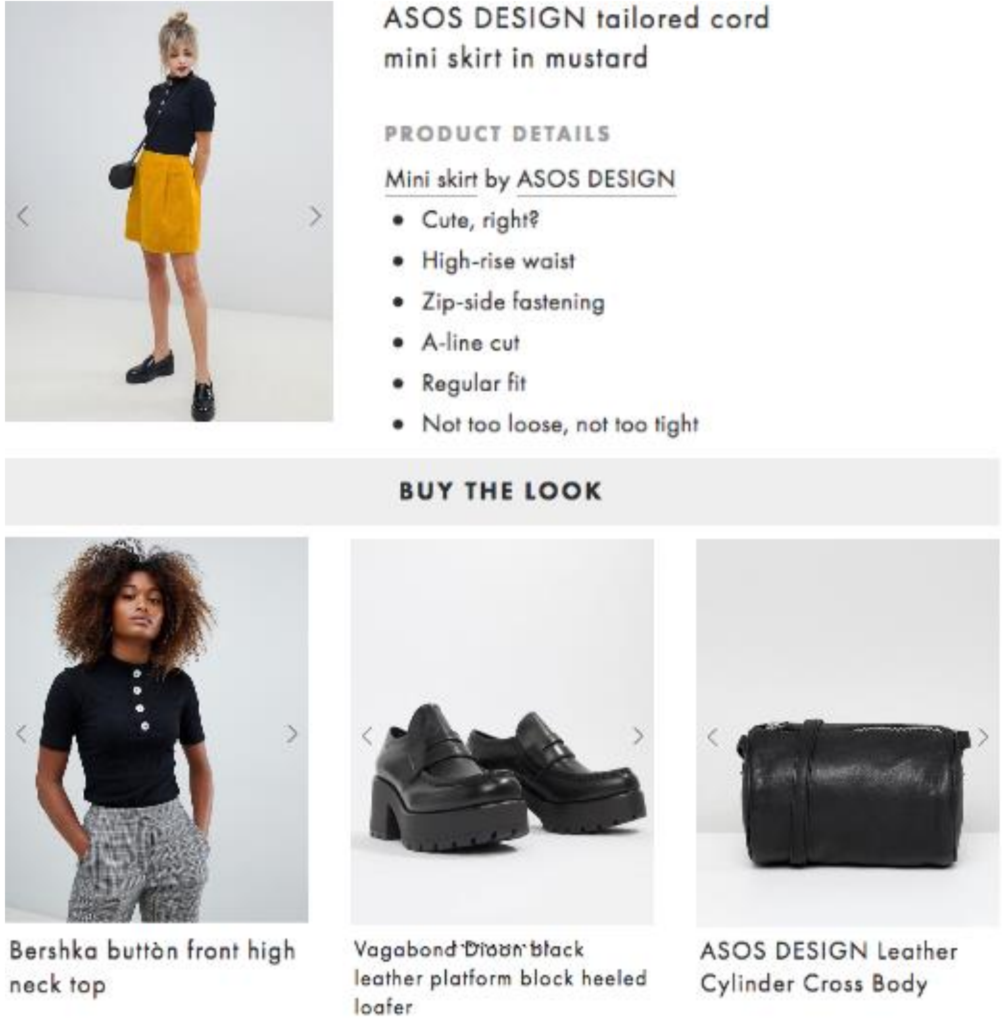

Source: [The ASOS Tech Blog, 2020]

Personalization

- Data: 600k outfits ("Buy The Look") curated by ASOS stylists (positive instance) + random changes to these outfits to create negative instance
  - ❑ What is a data instance? What is the target variable Y?
  - ❑ What are the input variables (features)?
- Classification model to distinguish between BTL and negative instances

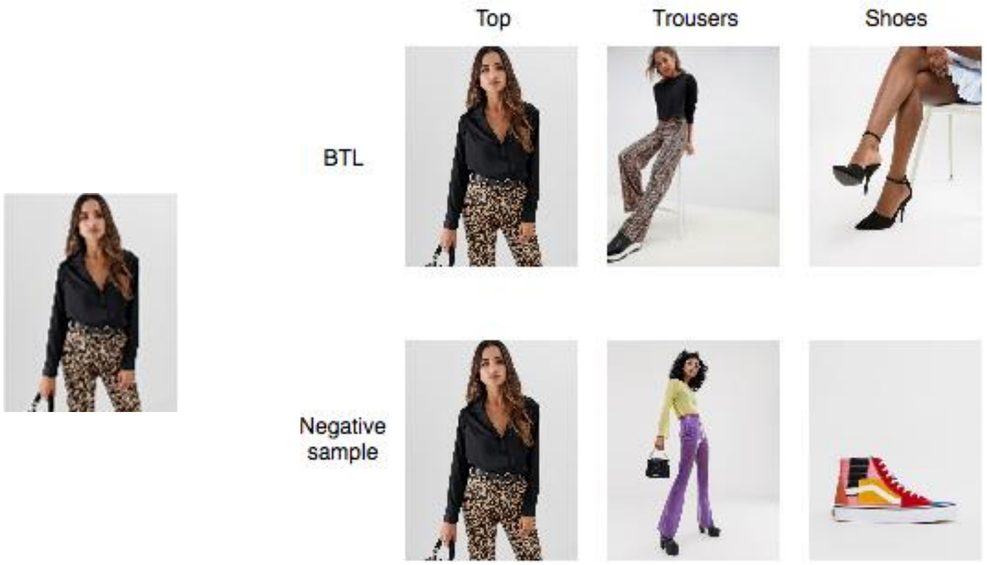

Source: [The ASOS Tech Blog, 2020]

Personalization

- Data: 600k outfits (“Buy The Look”) curated by ASOS stylists (positive instance) + random changes to these outfits to create negative instance
  - ❑ What is a data instance? What is the target variable Y?
  - ❑ What are the input variables (features)?
- Classification model to distinguish between BTL and negative instances

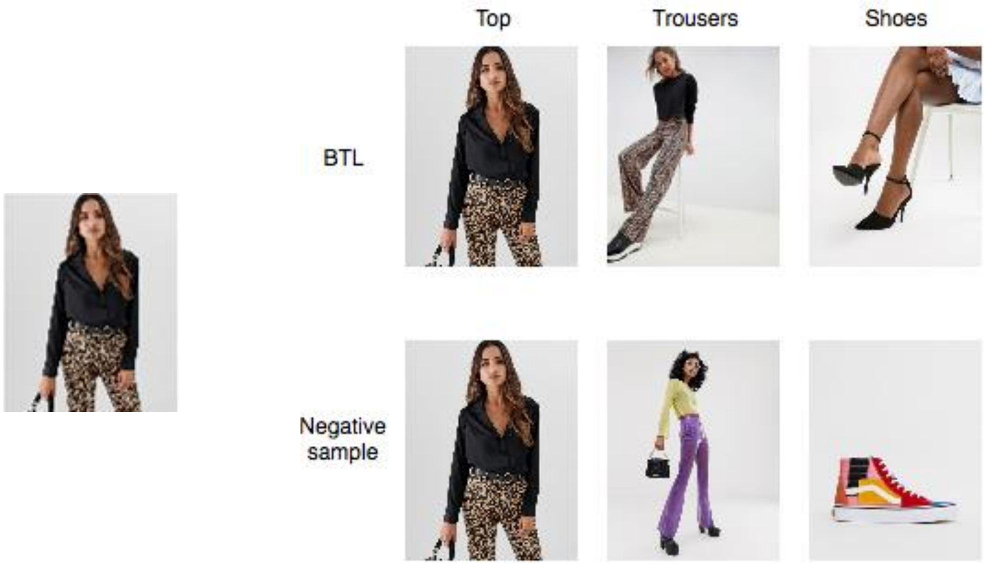

Target value: Buy The Look (Y=1)

Target value: Bad outfit (Y=0)

Source: [The ASOS Tech Blog, 2020]

Personalization

Personalization

Item embeddings represent an item (e.g., T-shirt)

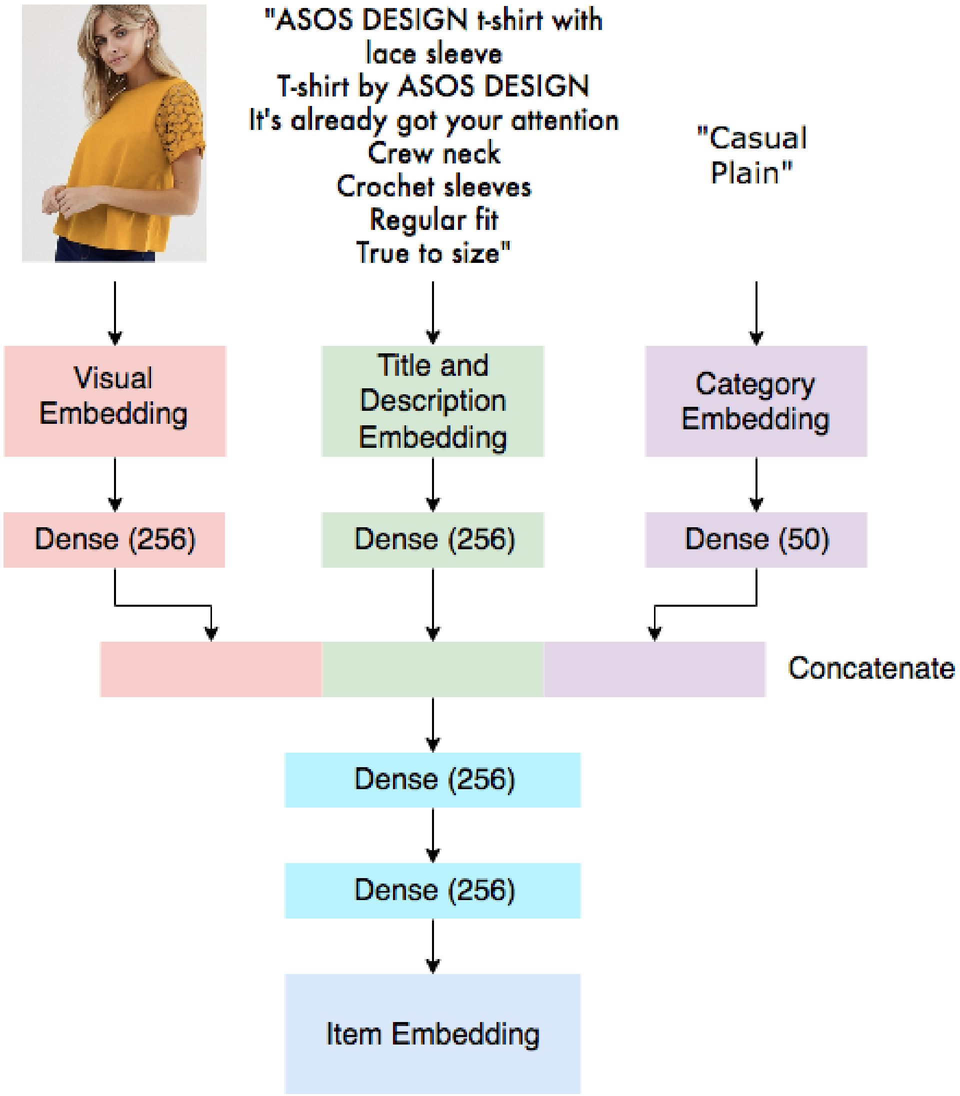

Source: [The ASOS Tech Blog, 2020]

Personalization

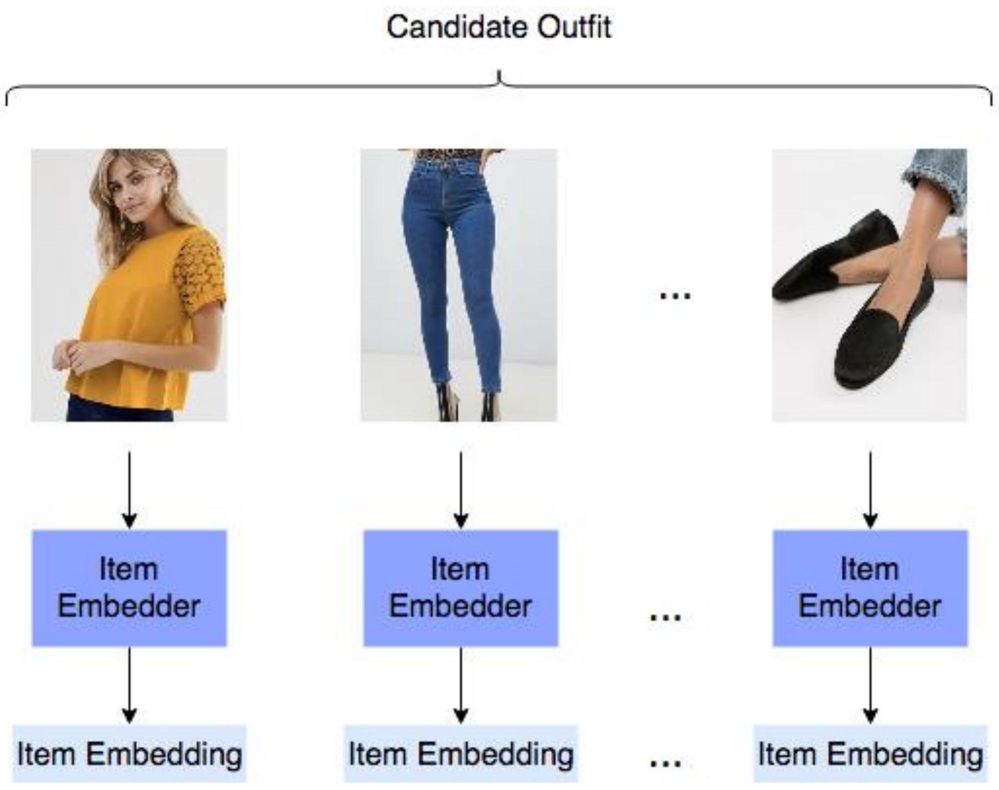

Source: [The ASOS Tech Blog, 2020]

Personalization

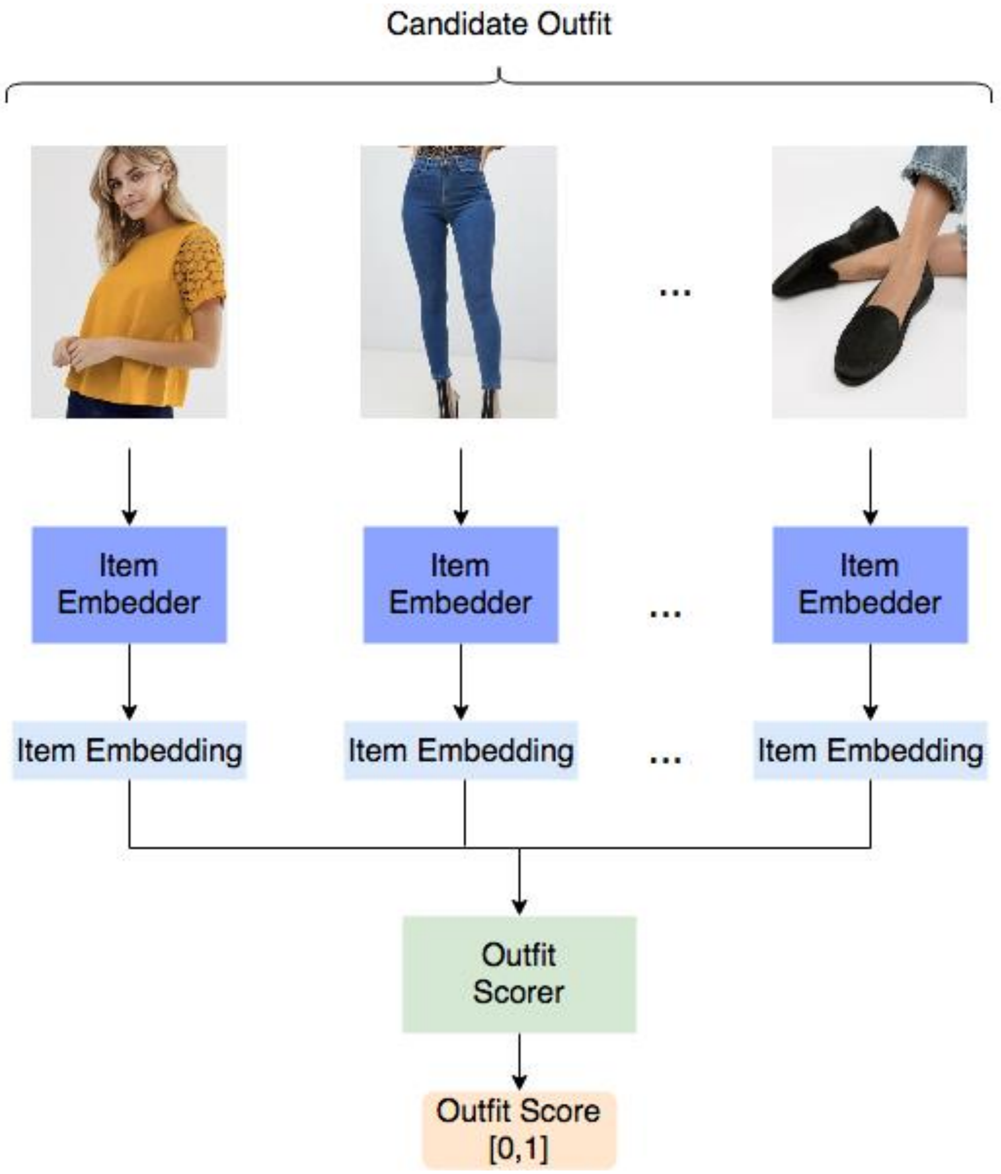

Source: [The ASOS Tech Blog, 2020](#)

Other applications

- Supply chain optimisation (inventory management, sales prediction, workforce optimization, delivery time & location, ...)
- HR analytics (automated resume screening, behavioral assessments, ...)
- Customer care optimisation (intent modeling, frustration detection, natural language processing, chatbots, ...)
- AI-driven products & services (self-driving cars, fraud detection, smart city applications, e-health, credit scoring, ...)
- …

4. Ethical challenges

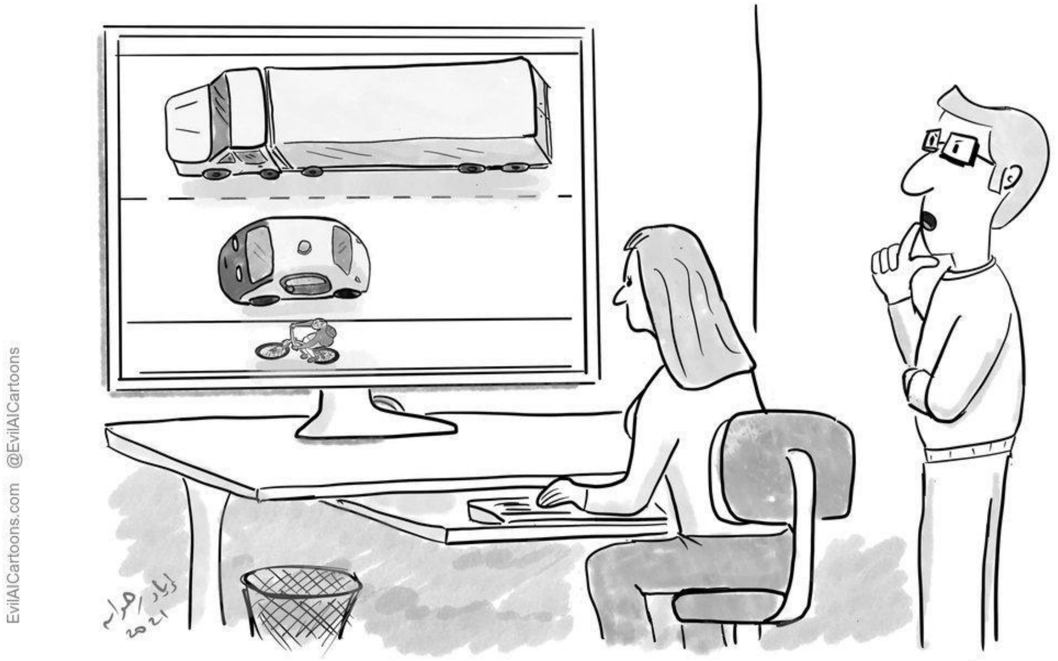

"If this happens, let's program the car to move 2 inches away from the truck, to reduce risk of collision. I don't like cyclists anyway!"

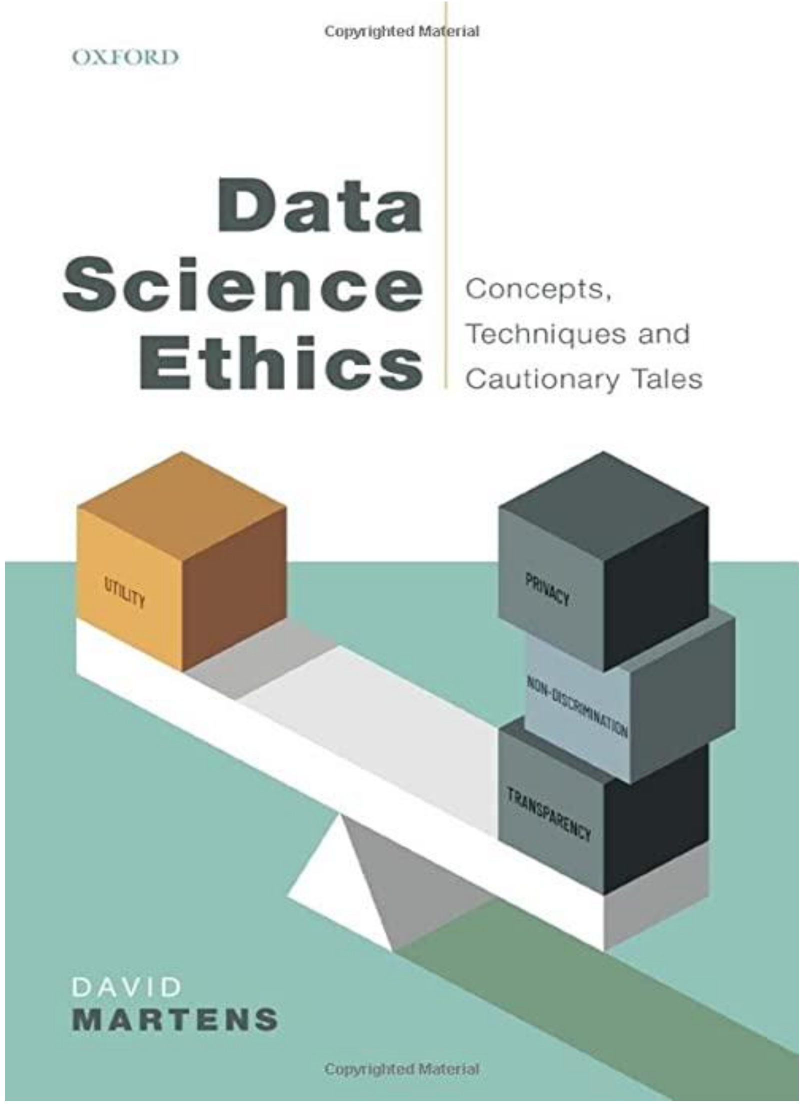

Privacy and transparency

"Hey! You're having a baby!" **Target**

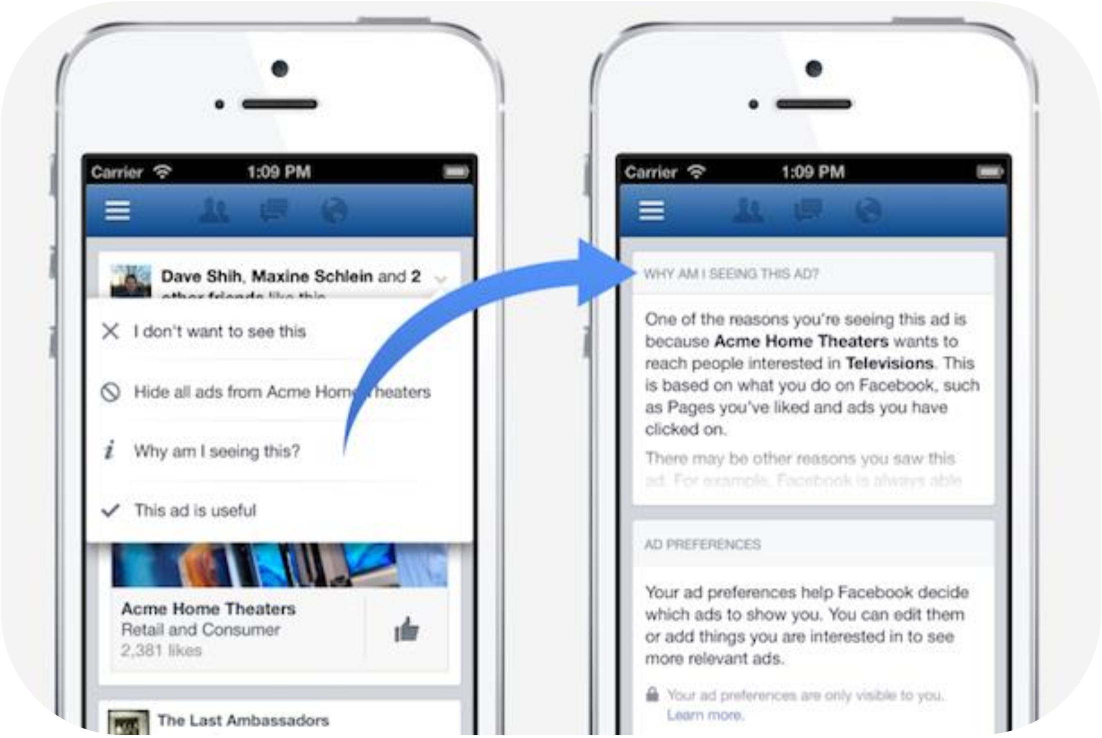

"Why am I seeing this ad?" **Facebook**

Fairness

[Racist soap dispenser]

[Racist Google Photos]

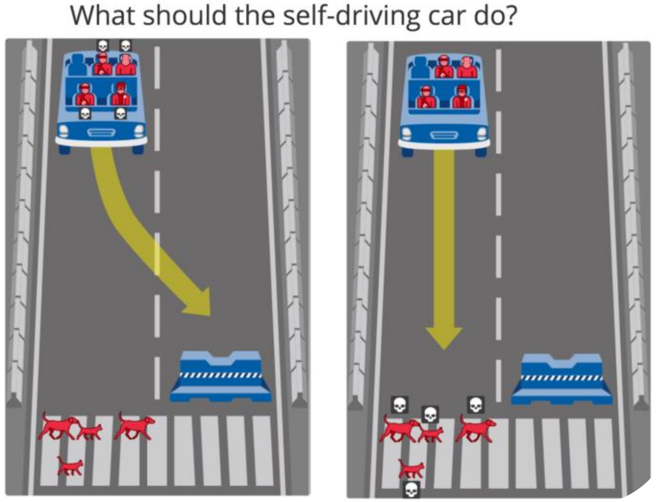

[MIT's Moral Machine]

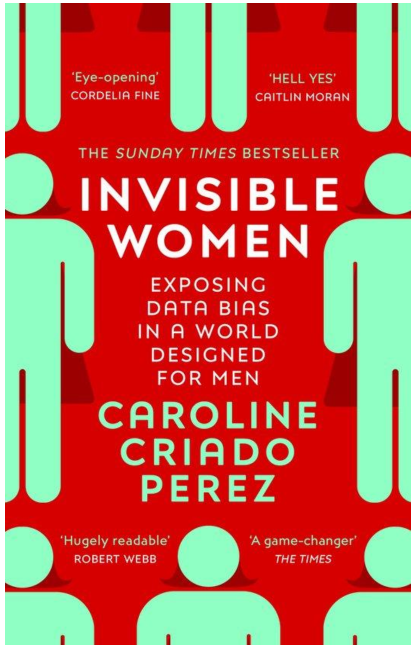

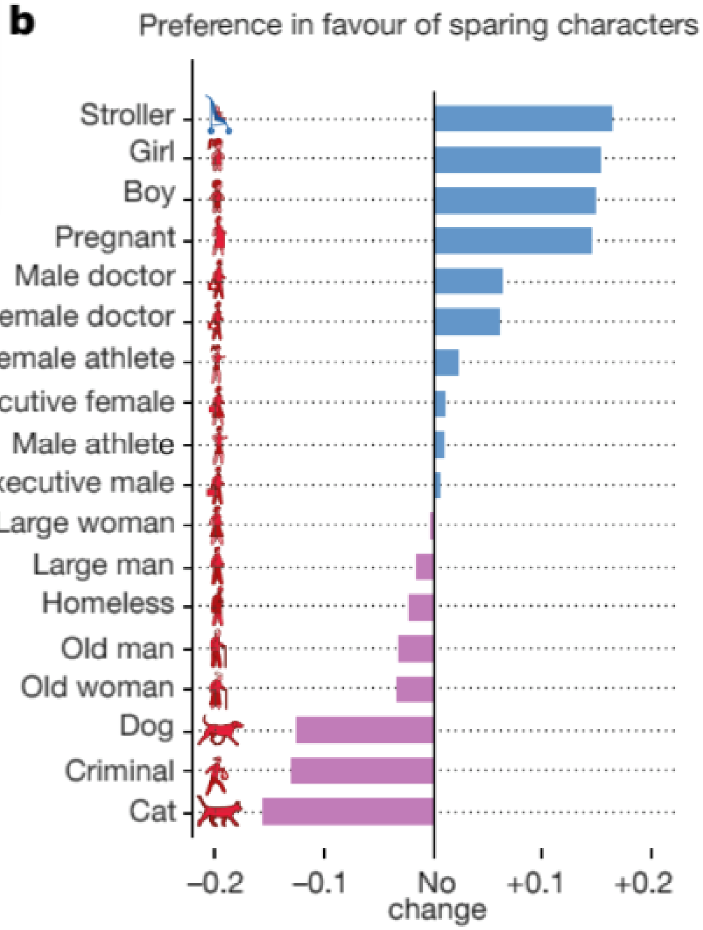

Explainability to avoid black-box models

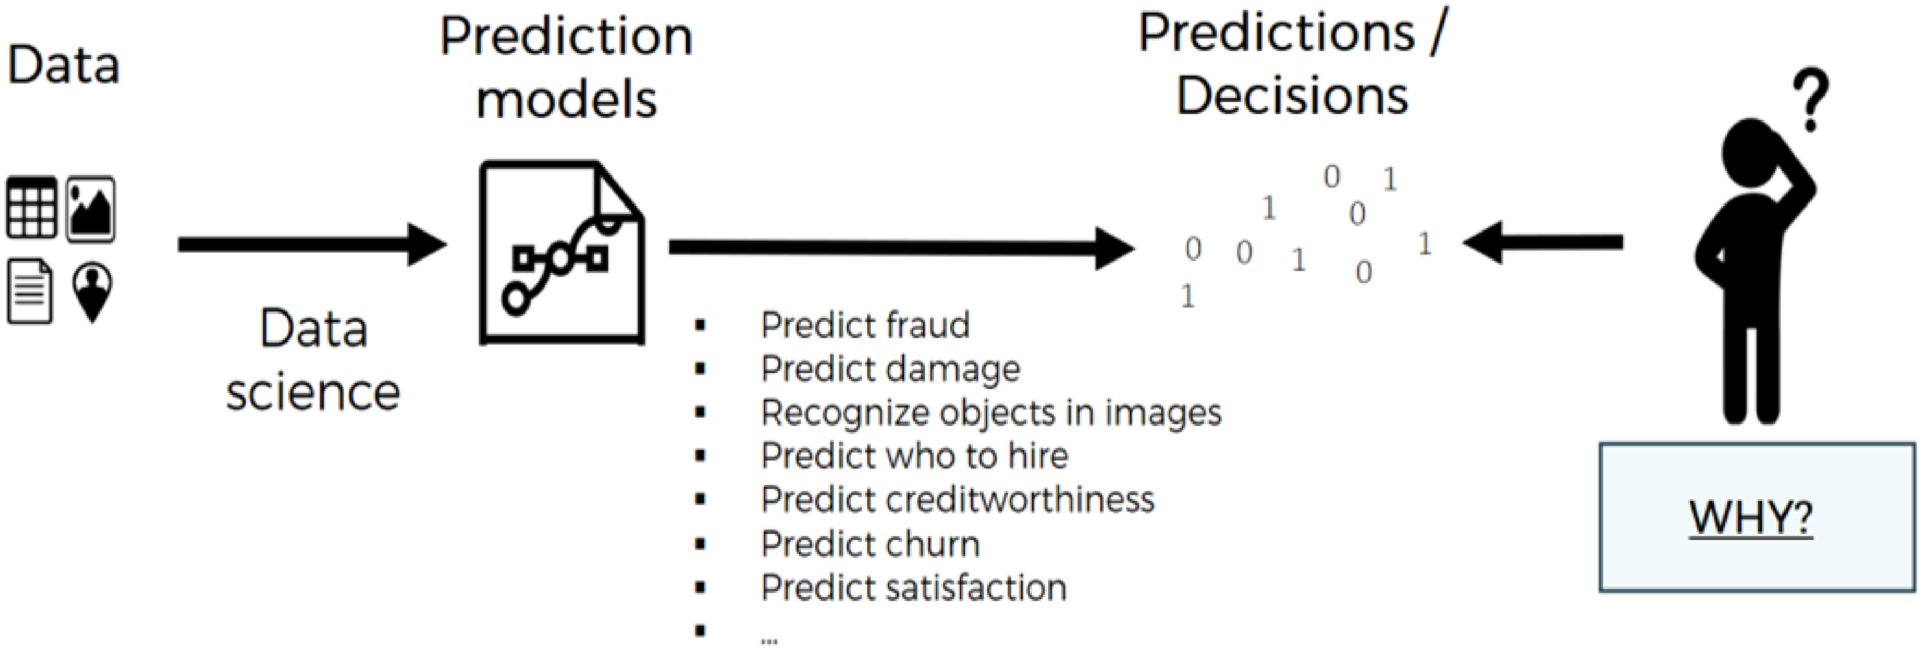

Source: [Data Science Ethics (Martens, 2022)]

What have you learned today?

You learned that ..

.. data mining is the automated extraction of a pattern from data.

.. data mining tasks vary in how and which pattern they extract from data.

.. the process of learning models from historical data to predict a (target) variable of interest is based on finding informative variables. The process includes a.o. preprocessing, modeling, and evaluation.

.. more data & processing power create opportunities to increase efficiency and/or accuracy of business functions and processes.

.. data science applications bring new ethical challenges related to data privacy, transparency, fairness, and explainability.

Thank you!
Questions?

yanou.ramon@gmail.com

www.linkedin.com/in/yanouramon

yramon.github.io