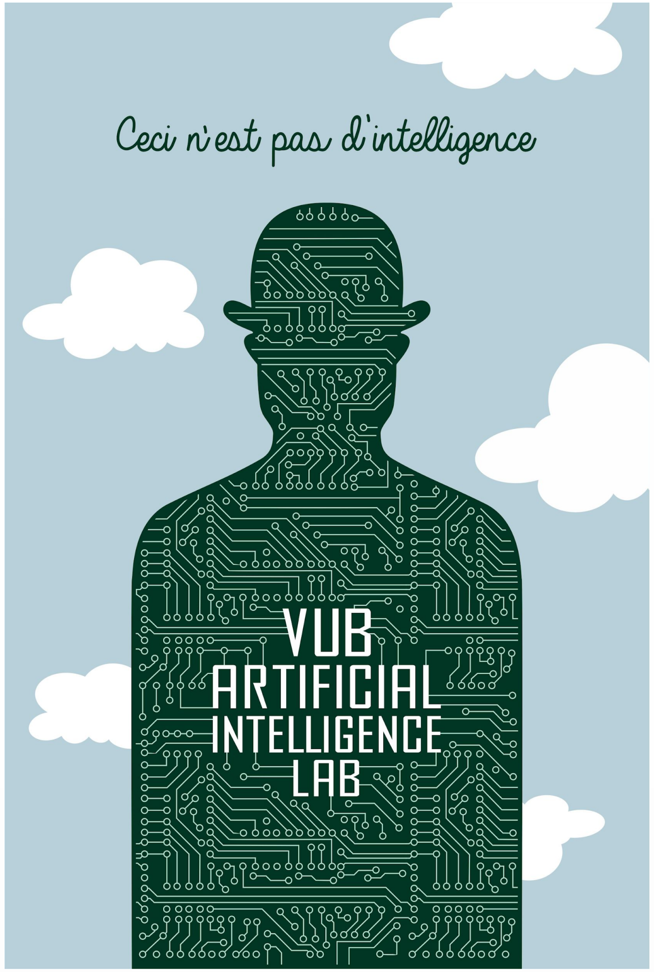

Artificial Intelligence

A broader perspective

Current Trends in AI

Johan Loeckx

28/02/2020

jloeckx@ai.vub.ac.be

Q#1
Give examples of "AI in use"


BBRFC Celtic
Rugby Club
R21
R21
R21
N206
8 min
650 m
10 min
750 m
Beobank
Le Quartier Latin
Etterbeek
Luigi's
Work
VUB Brussels Building B
Vrije
Universiteit
Brussel
Vesalius College vzw/asbl
Vrije Universiteit
Brussel Building ROT2
VUB - Vrije Universiteit
Brussel Building ROT3
Building Q
VUB - Vrije Universiteit
Brussel Building ROT4
VUB - Vrije Universiteit
Brussel Building F
Vrije Universiteit
Brussel Building G
Ing VUB
Avenue de la Cou

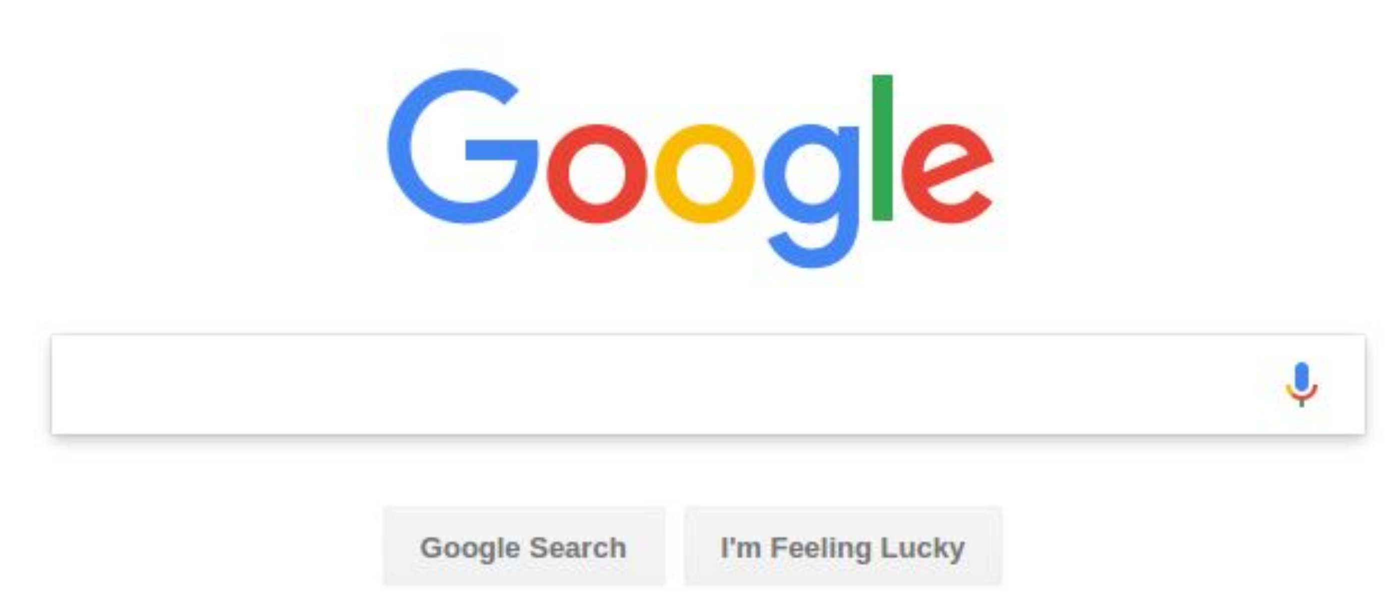
Google
Google Search
I'm Feeling Lucky

VISA

		3	49	2	147	6		
9			3	47	5			1
		1	8	79	6	4		
		8	1		2	9		
7								8
		6	7		8	2		
		2	6		9	5		
8			2		3			9
		5		1		3		

LOGIC

(e.g. model expansion)

Convolutionary neural networks

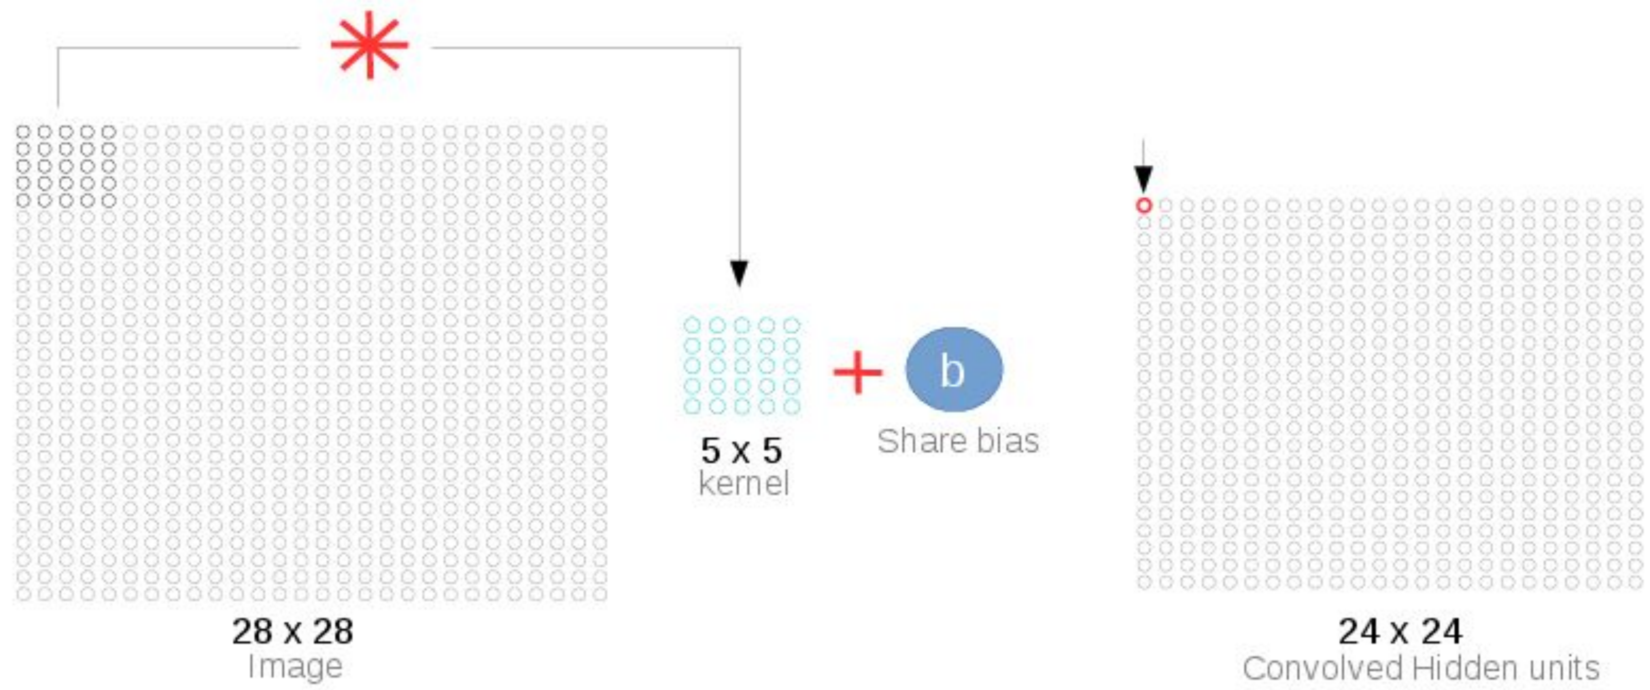

Convolutional neural network

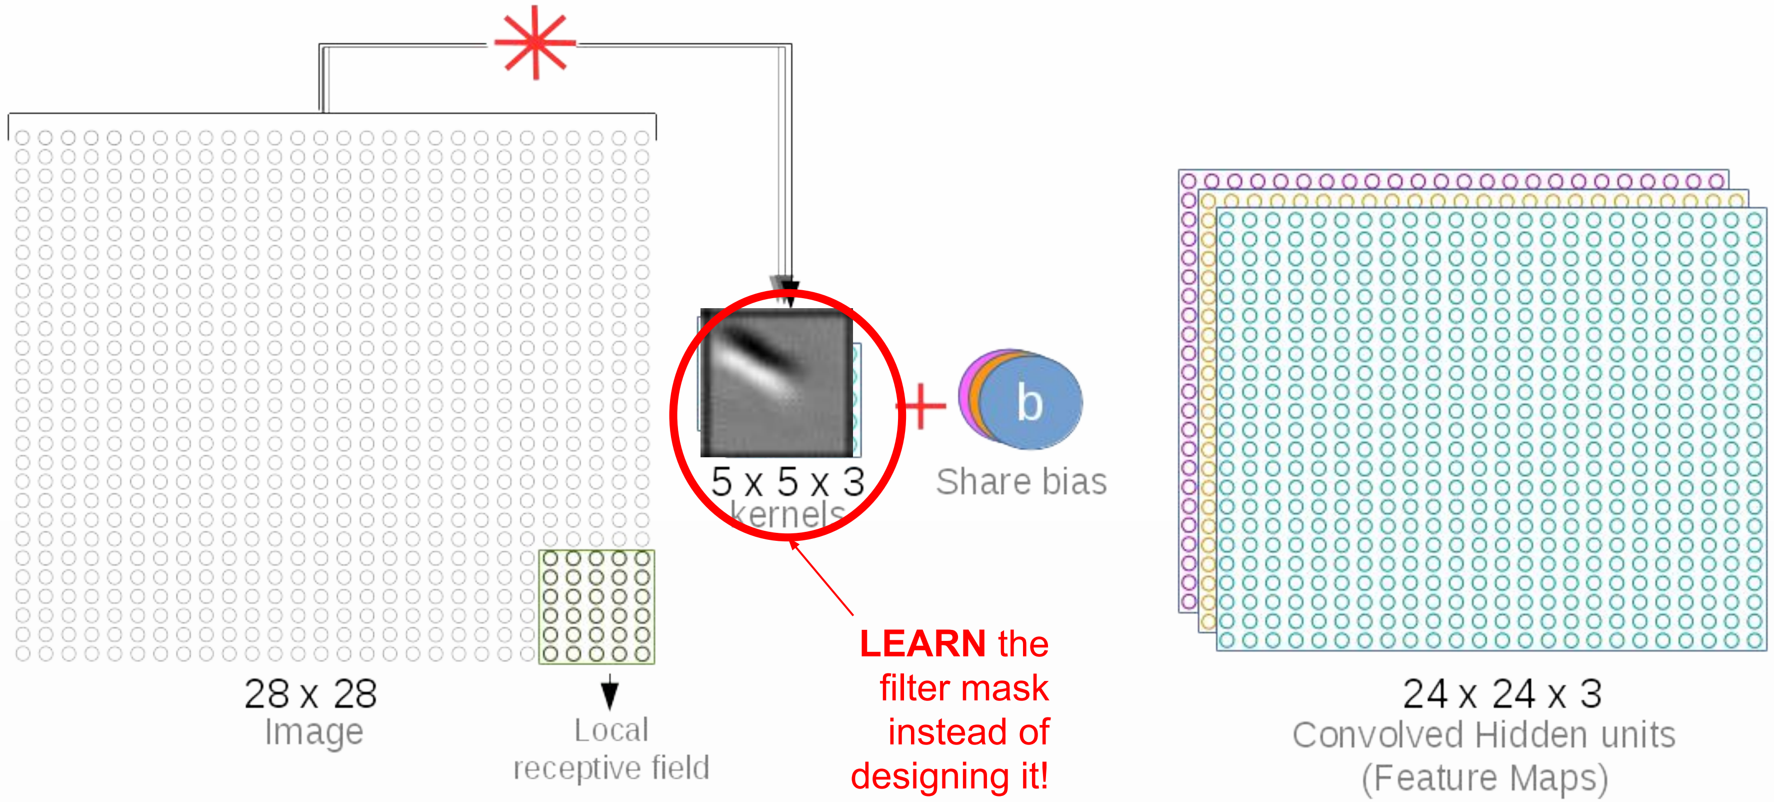

LEARN the filter mask instead of designing it!

Recommendation

- Given a sparse matrix of (user,item) preferences
- Predict new preferences

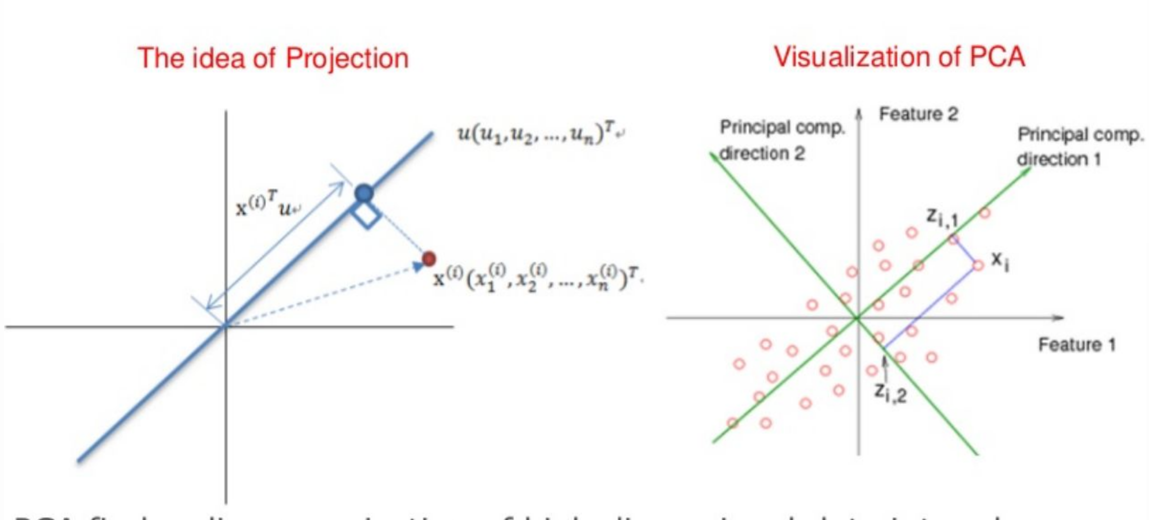

PCA finds a linear projection of high dimensional data into a lower dimensional subspace.

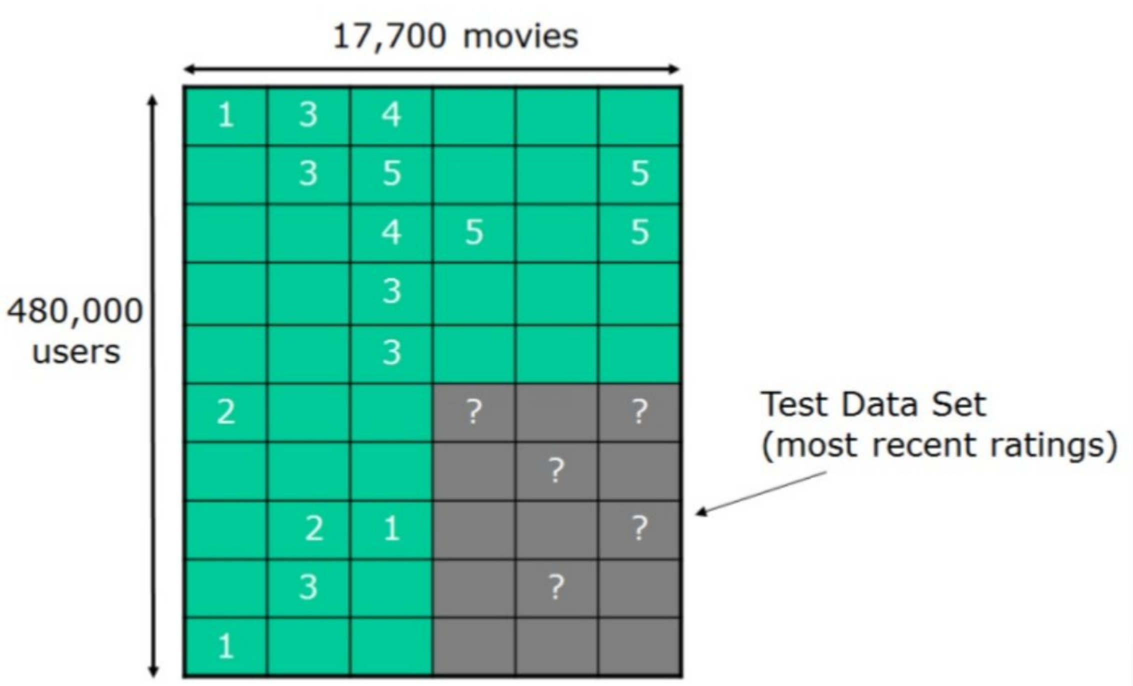

Object-oriented programming

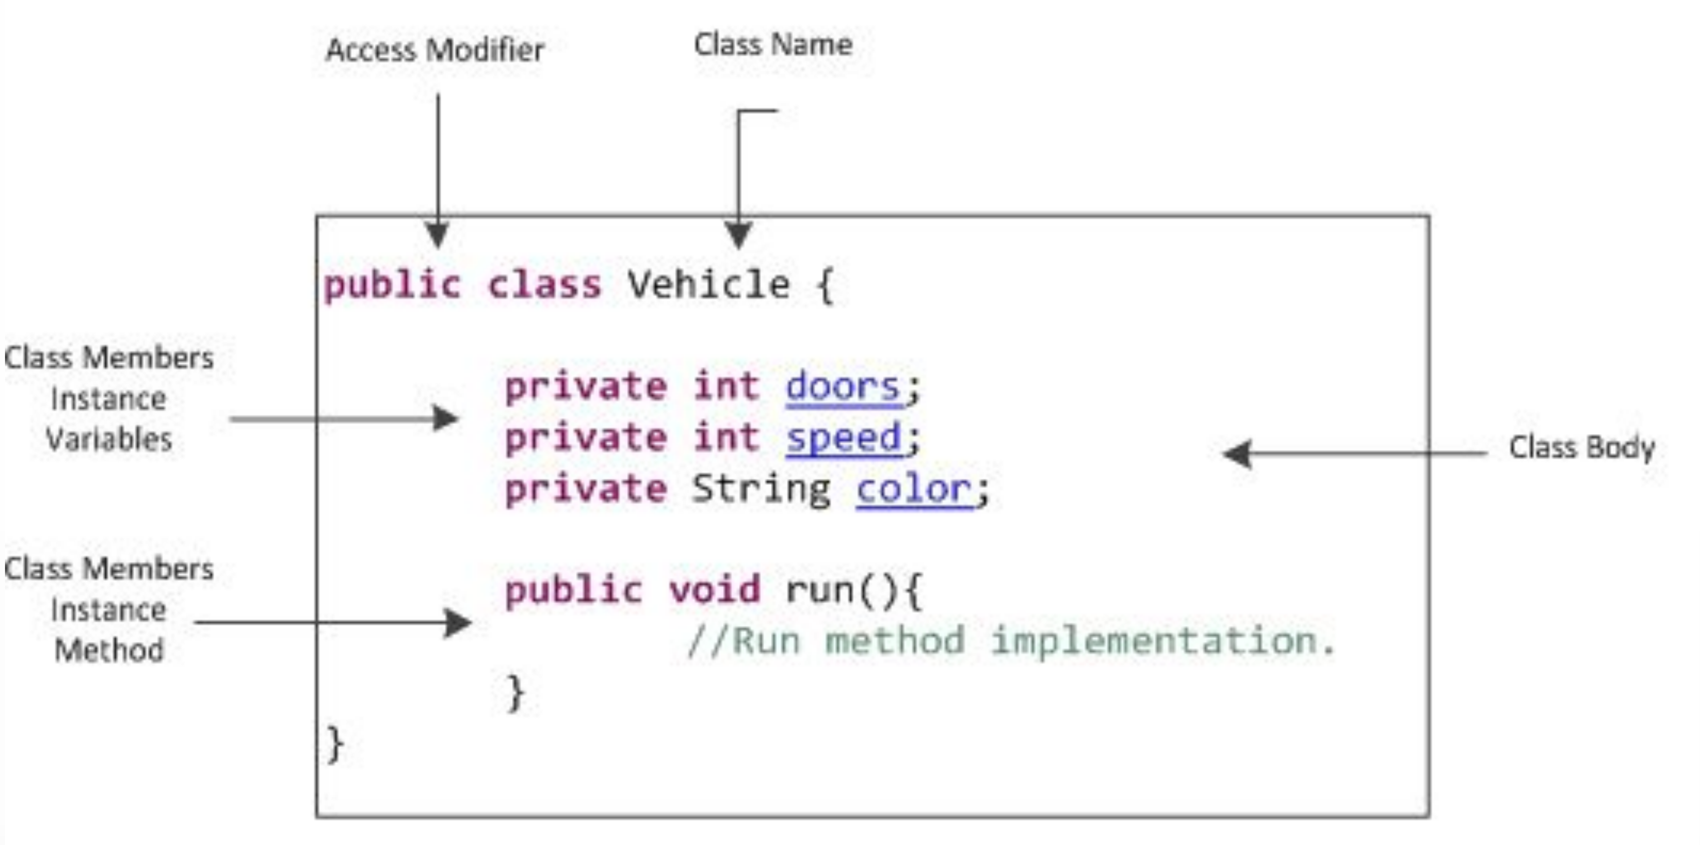

Natural language

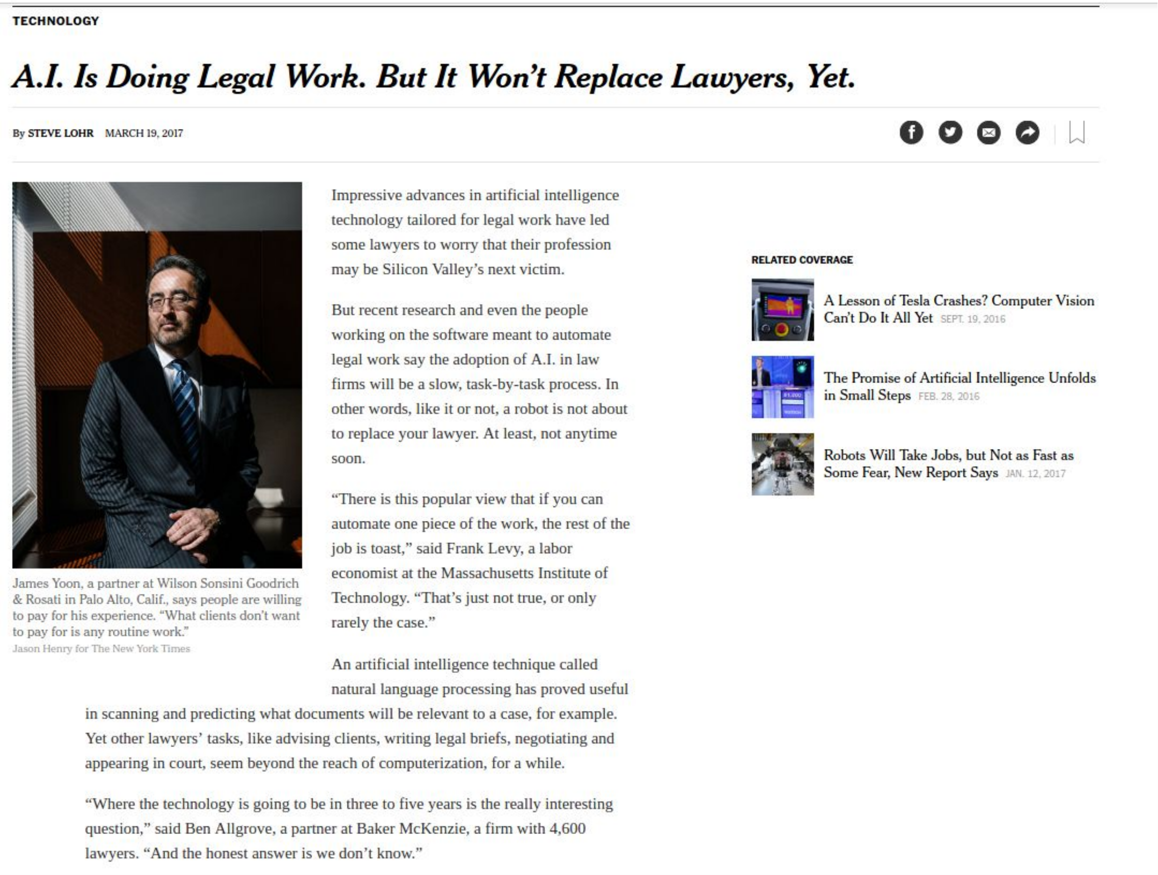

TECHNOLOGY

A.I. Is Doing Legal Work. But It Won't Replace Lawyers, Yet.

By STEVE LOHR MARCH 19, 2017

James Yoon, a partner at Wilson Sonsini Goodrich & Rosati in Palo Alto, Calif., says people are willing to pay for his experience. "What clients don't want to pay for is any routine work."
Jason Henry for The New York Times

Impressive advances in artificial intelligence technology tailored for legal work have led some lawyers to worry that their profession may be Silicon Valley's next victim.

But recent research and even the people working on the software meant to automate legal work say the adoption of A.I. in law firms will be a slow, task-by-task process. In other words, like it or not, a robot is not about to replace your lawyer. At least, not anytime soon.

"There is this popular view that if you can automate one piece of the work, the rest of the job is toast," said Frank Levy, a labor economist at the Massachusetts Institute of Technology. "That's just not true, or only rarely the case."

An artificial intelligence technique called natural language processing has proved useful in scanning and predicting what documents will be relevant to a case, for example. Yet other lawyers' tasks, like advising clients, writing legal briefs, negotiating and appearing in court, seem beyond the reach of computerization, for a while.

"Where the technology is going to be in three to five years is the really interesting question," said Ben Allgrove, a partner at Baker McKenzie, a firm with 4,600 lawyers. "And the honest answer is we don't know."

RELATED COVERAGE

A Lesson of Tesla Crashes? Computer Vision Can't Do It All Yet SEPT. 19, 2016

The Promise of Artificial Intelligence Unfolds in Small Steps FEB. 28, 2016

Robots Will Take Jobs, but Not as Fast as Some Fear, New Report Says JAN. 12, 2017

- Sentiment analysis
- Summarization
- Classification & routing
- Named entities, events

ARTIFICIAL INTELLIGENCE RESEARCH GROUP

@aibrussels

Q#2
What is AI?

AI studies…

1. the nature and mechanisms of **intelligence**
2. Using **formal methods**, an
3. Attempts to **reconstruct** it.

in **principle**

Playing chess...

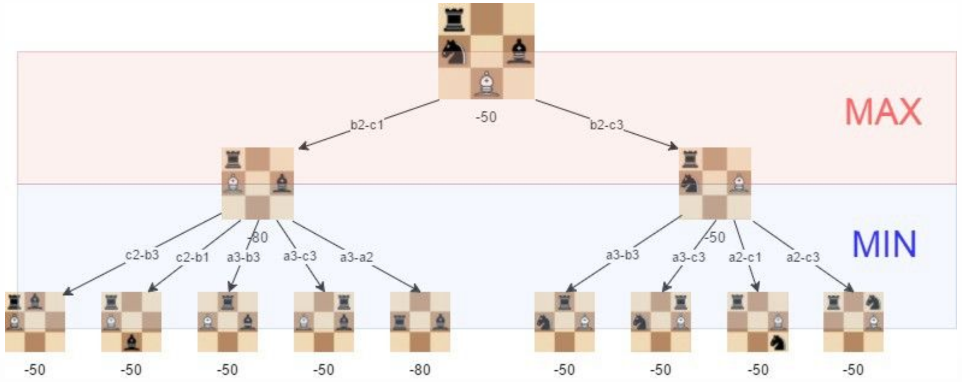

In principle: learn **flying**, not the **bird**.

- Do not copy the flapping wings
- Learn how a bird uses physics to fly

The question of whether a computer can think is no more interesting than the question of whether a submarine can swim -- Edsger W. Dijkstra

It does **not serve a single function**

- It's a **mishmash** of techniques, methods & ideas
- It's always **at the front** of computer science and is in continuous **change**
- It embodies the **quest to express meaning** in machines: a number can represent anything!

3 waves of artificial intelligence

Waves of artificial intelligence

1. Manual coding
   → if-then-rules, ad-hoc
2. **Formalizing** intelligence
   → reformulate in a structured way (e.g. search)
3. Learning from **data**
   → statistical patterns
4. Learning from **interaction**
   → learn what *decisions* to take from rewards

learn from human knowledge, not data

Role of the human when designing AI

1. Manual ad-hoc coding (not really called AI)

Role of the human when designing AI

2. **Symbolic** AI: learning from encoded **knowledge**

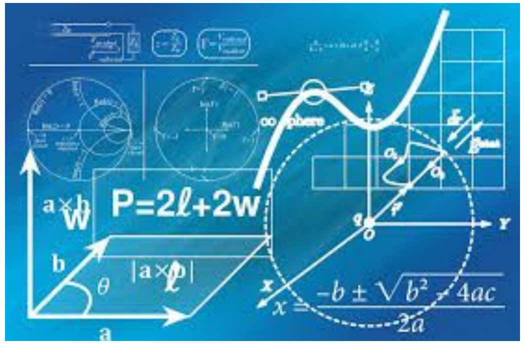

formalism!!

build upon existing implementations

Role of the human when designing AI

3. Learning from **data**

what **data** to collect? how
to represent reality?

what **task** to solve?

time	n.risk	n.event	survival	std.err	lower 95% CI	upper 95% CI
5	23	2	0.913	0.0588	0.8049	1
8	21	2	0.8261	0.079	0.6848	0.996
9	19	1	0.7826	0.086	0.631	0.971
12	18	1	0.7391	0.0916	0.5798	0.942
13	17	1	0.6957	0.0959	0.5309	0.912
18	14	1	0.646	0.1011	0.4753	0.878
23	13	2	0.5466	0.1073	0.3721	0.803
27	11	1	0.4969	0.1084	0.324	0.762
30	9	1	0.4417	0.1095	0.2717	0.718
31	8	1	0.3865	0.1089	0.2225	0.671
33	7	1	0.3313	0.1064	0.1765	0.622
34	6	1	0.2761	0.102	0.1338	0.569
43	5	1	0.2208	0.0954	0.0947	0.515
45	4	1	0.1656	0.086	0.0598	0.458
48	2	1	0.0828	0.0727	0.0148	0.462

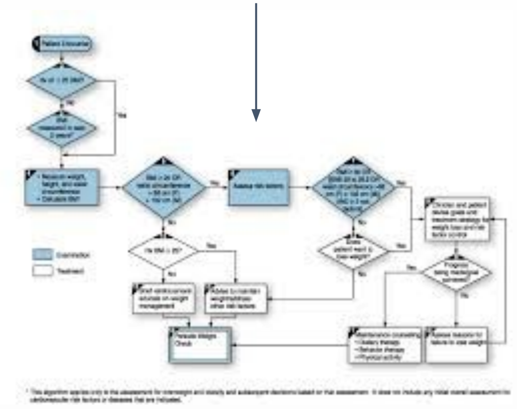

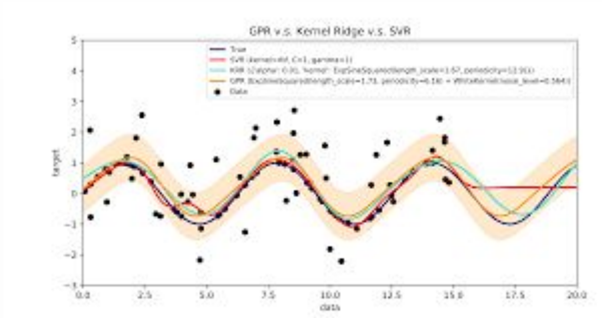

what **algorithm** &
preprocessing?

Role of the human when designing AI

4. Learning from **interaction** (reinforcement learning)

What does “success” mean?

actions

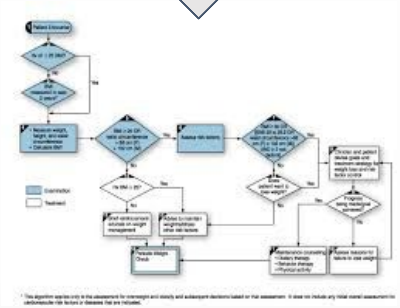

impact on

environment

TRENDS IN ML + ROLE HUMANS

from
knowledge

Symbolic AI,
learns to **reason**

Human creates **rules**

from
examples

Machine Learning,
learns a **mapping**

Humans annotate **data**

from
interaction

Reinforcement
Learning,
learns a **policy**

Humans design **success**

ARTIFICIAL
INTELLIGENCE
RESEARCH GROUP

But it is

Hard to structure.

From a **scope** point-of-view

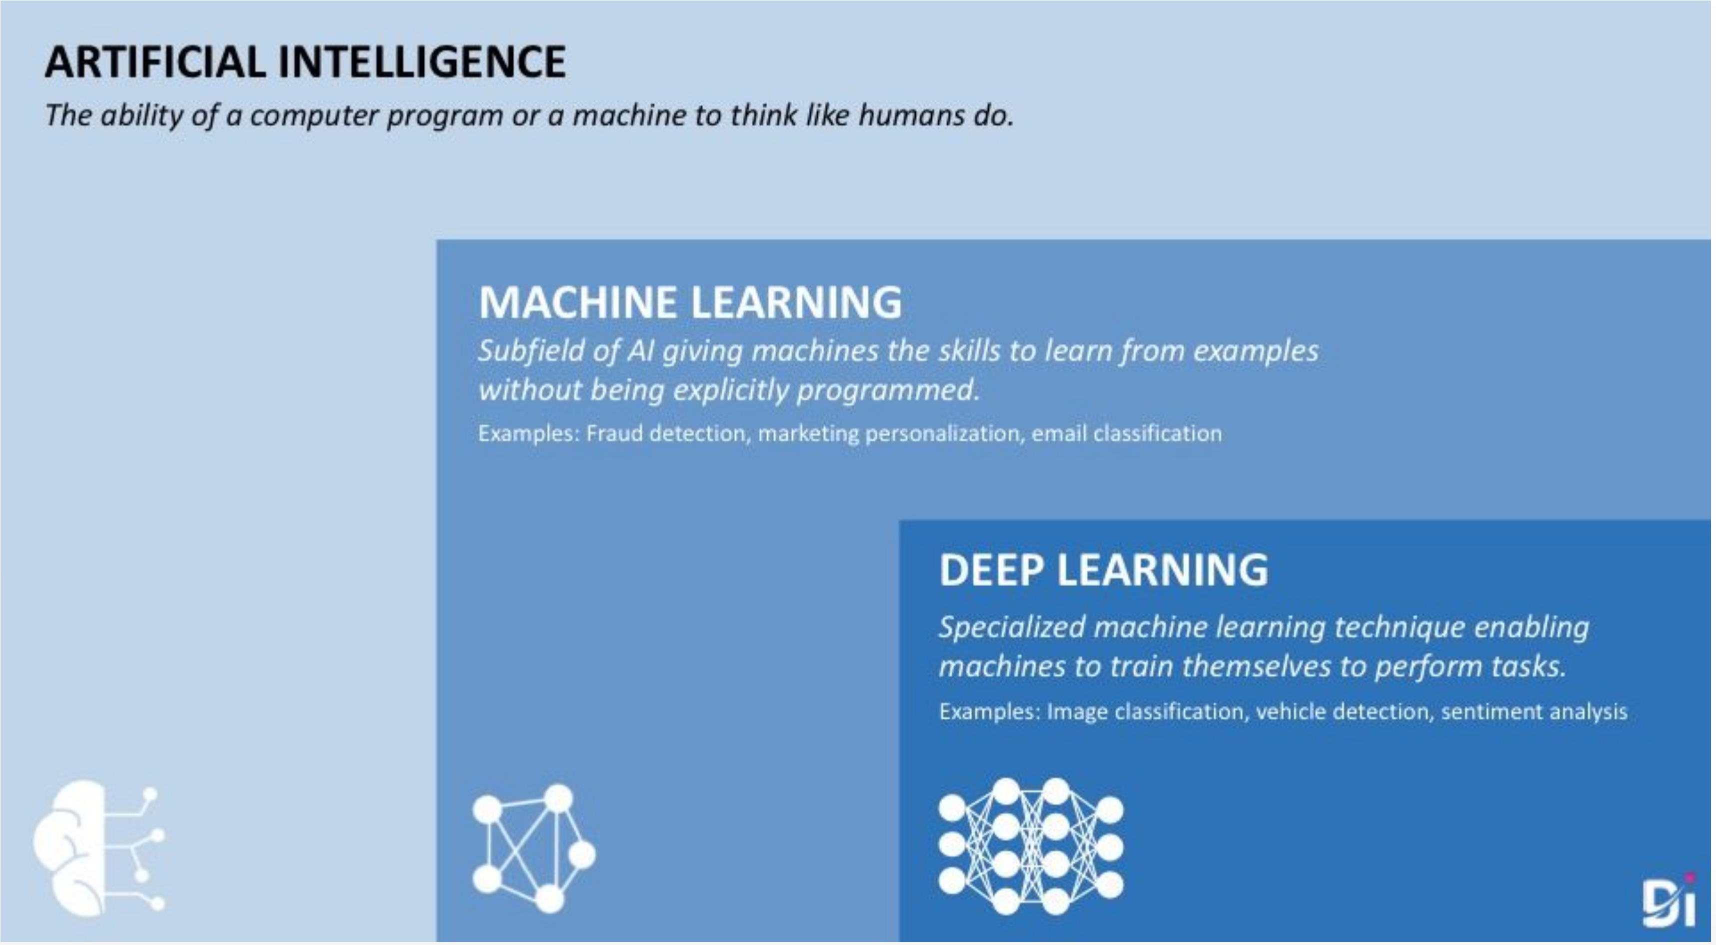

From an **algorithmic** point of view

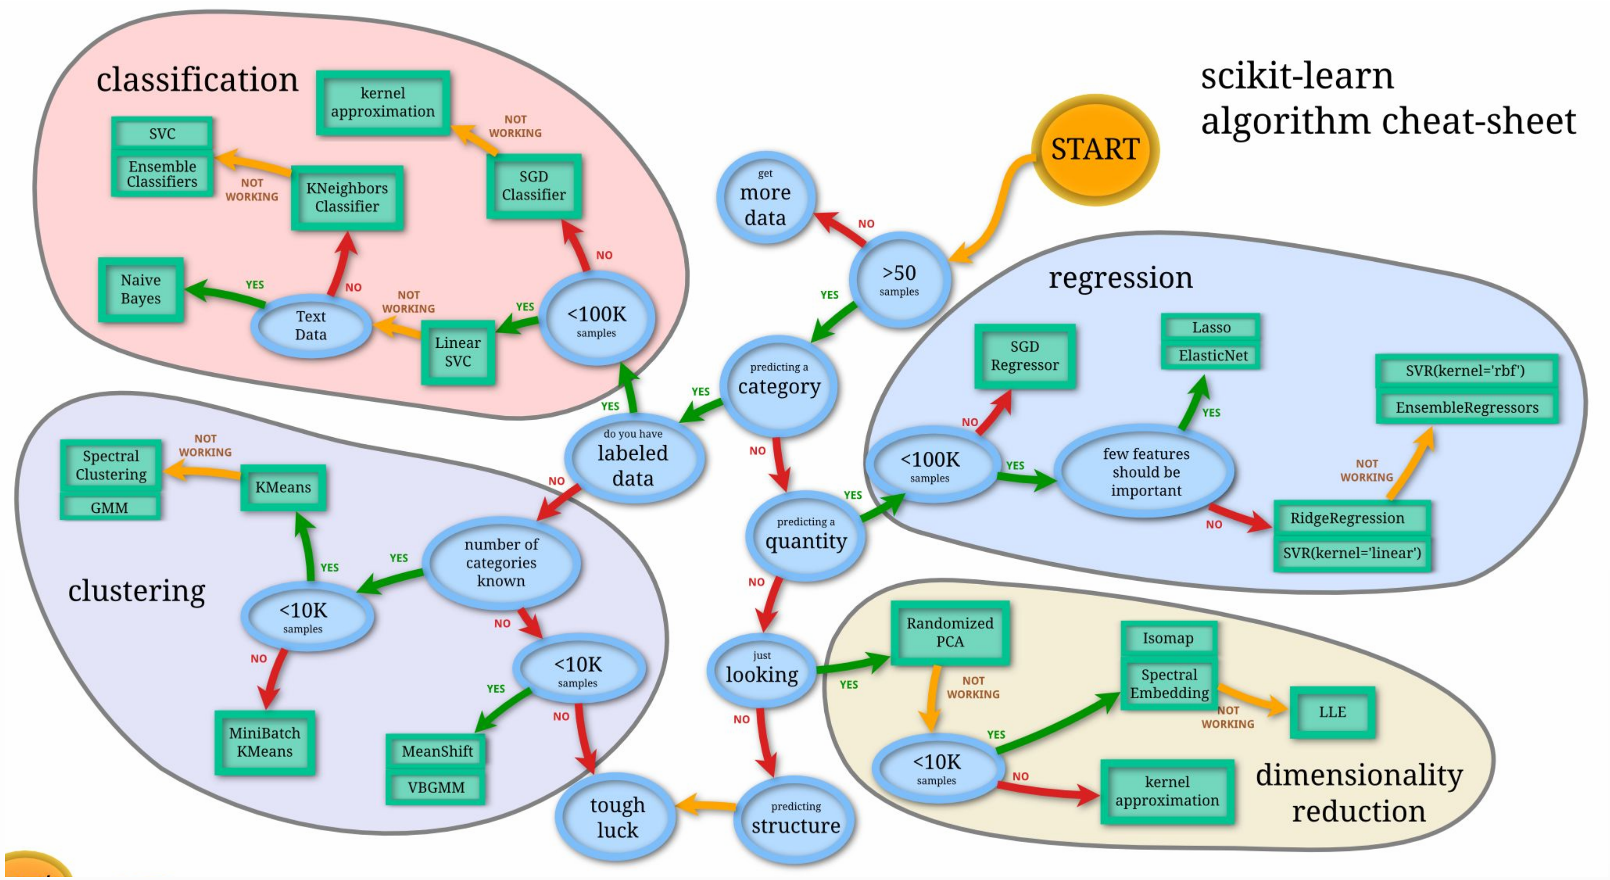

From an **outcome** point-of-view

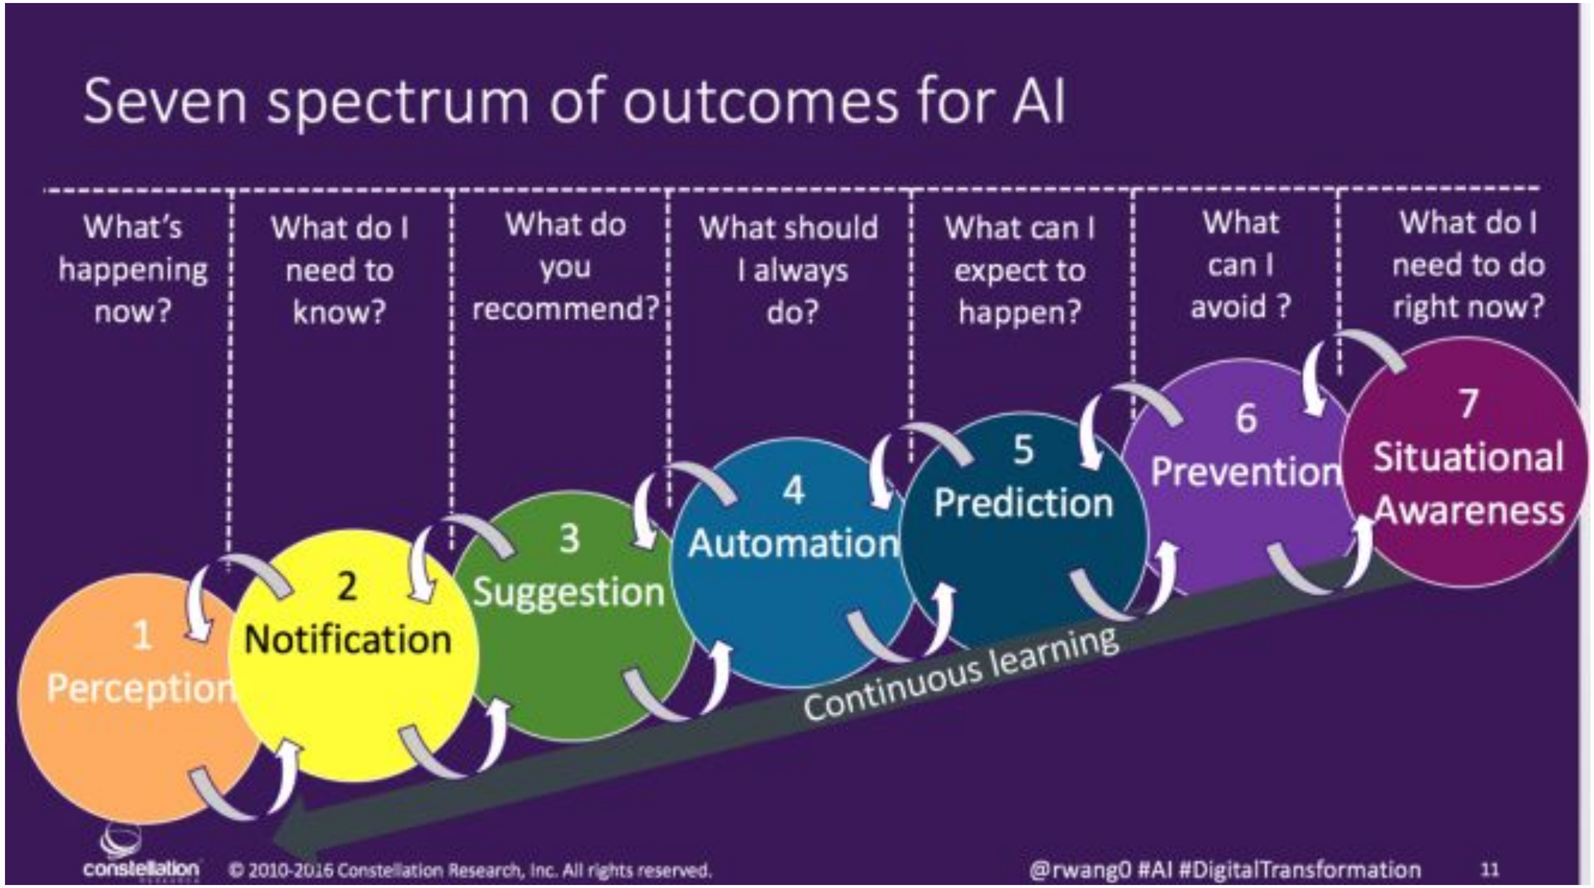

https://davidcummings.org/2016/12/28/seven-spectrum-of-outcomes-for-ai/

From an **application** point-of-view

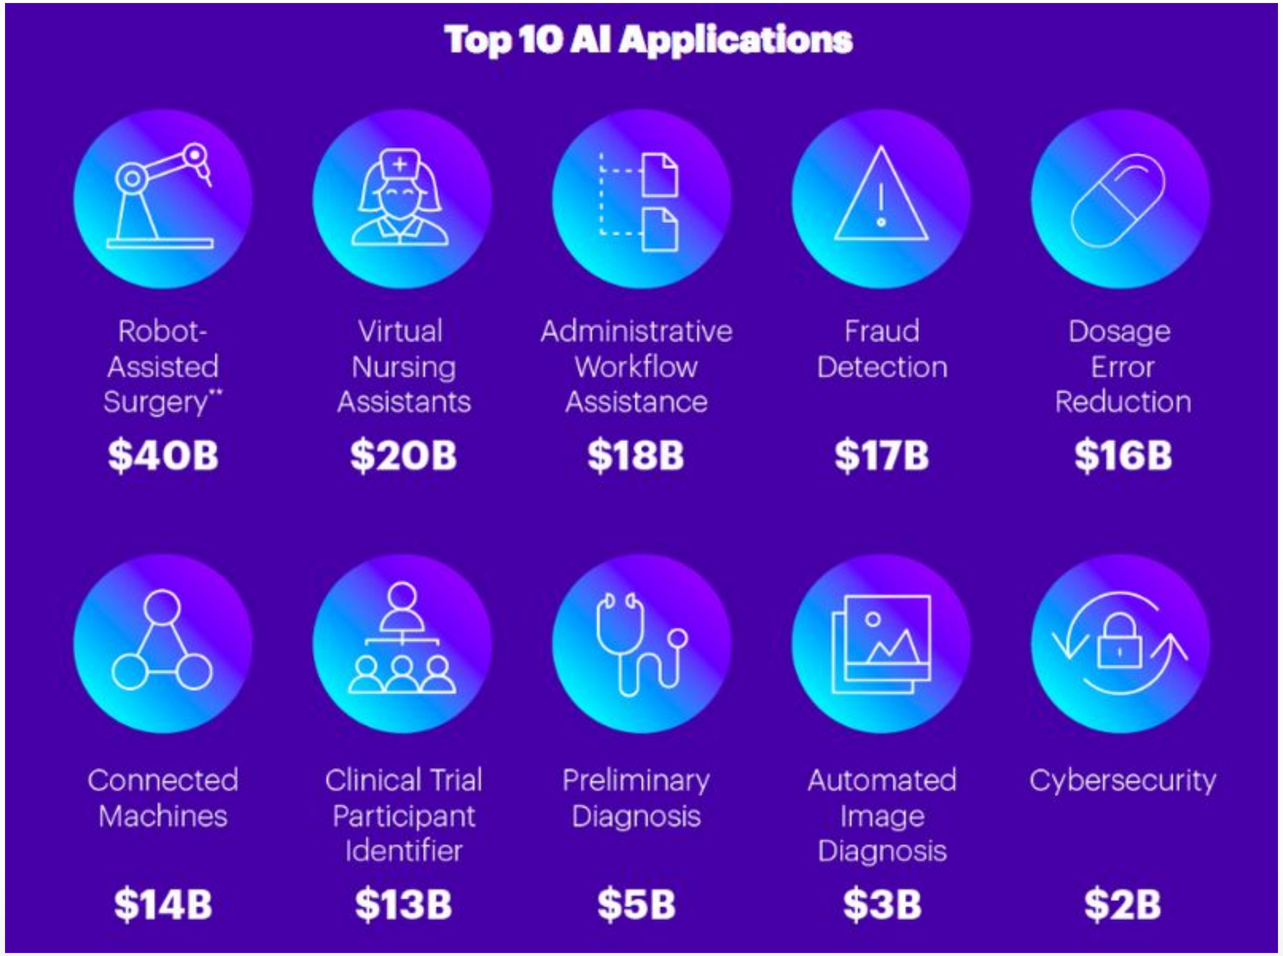

https://www.plugandplaytechcenter.com/resources/ai-healthcare/

From an **added-value** perspective

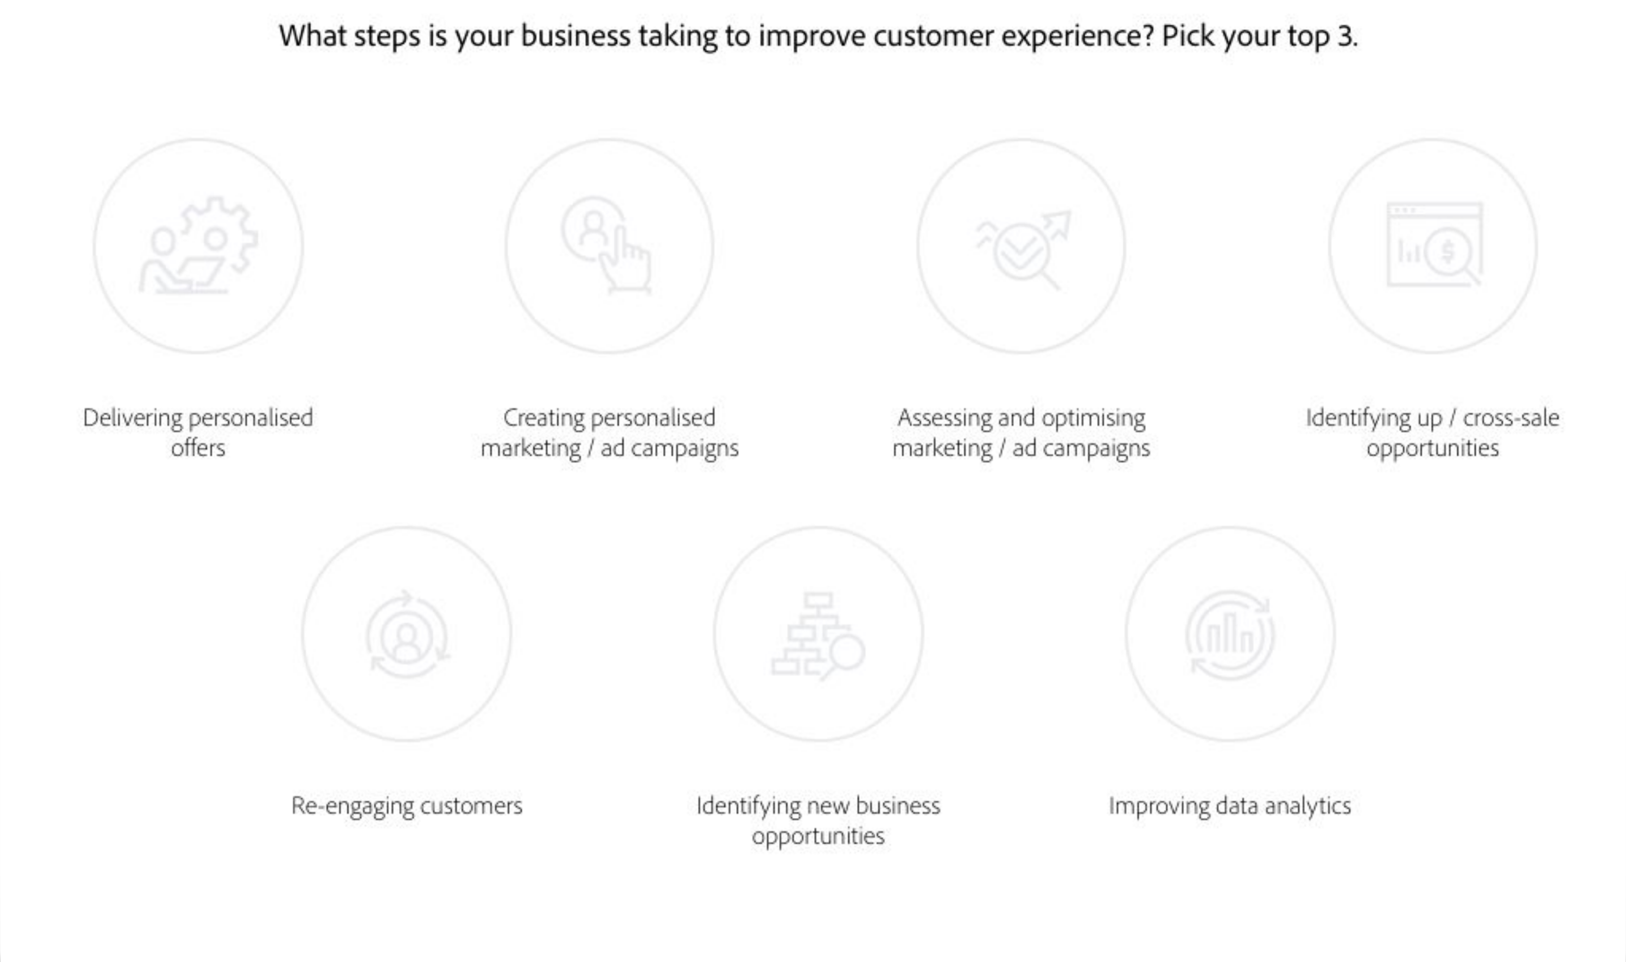

https://www.adobe.com/uk/artificial-intelligence/business-readiness-tool.html?sdid=55KD8VHP&mv=social&mv2=ownsoc

From **research field** point-of-view

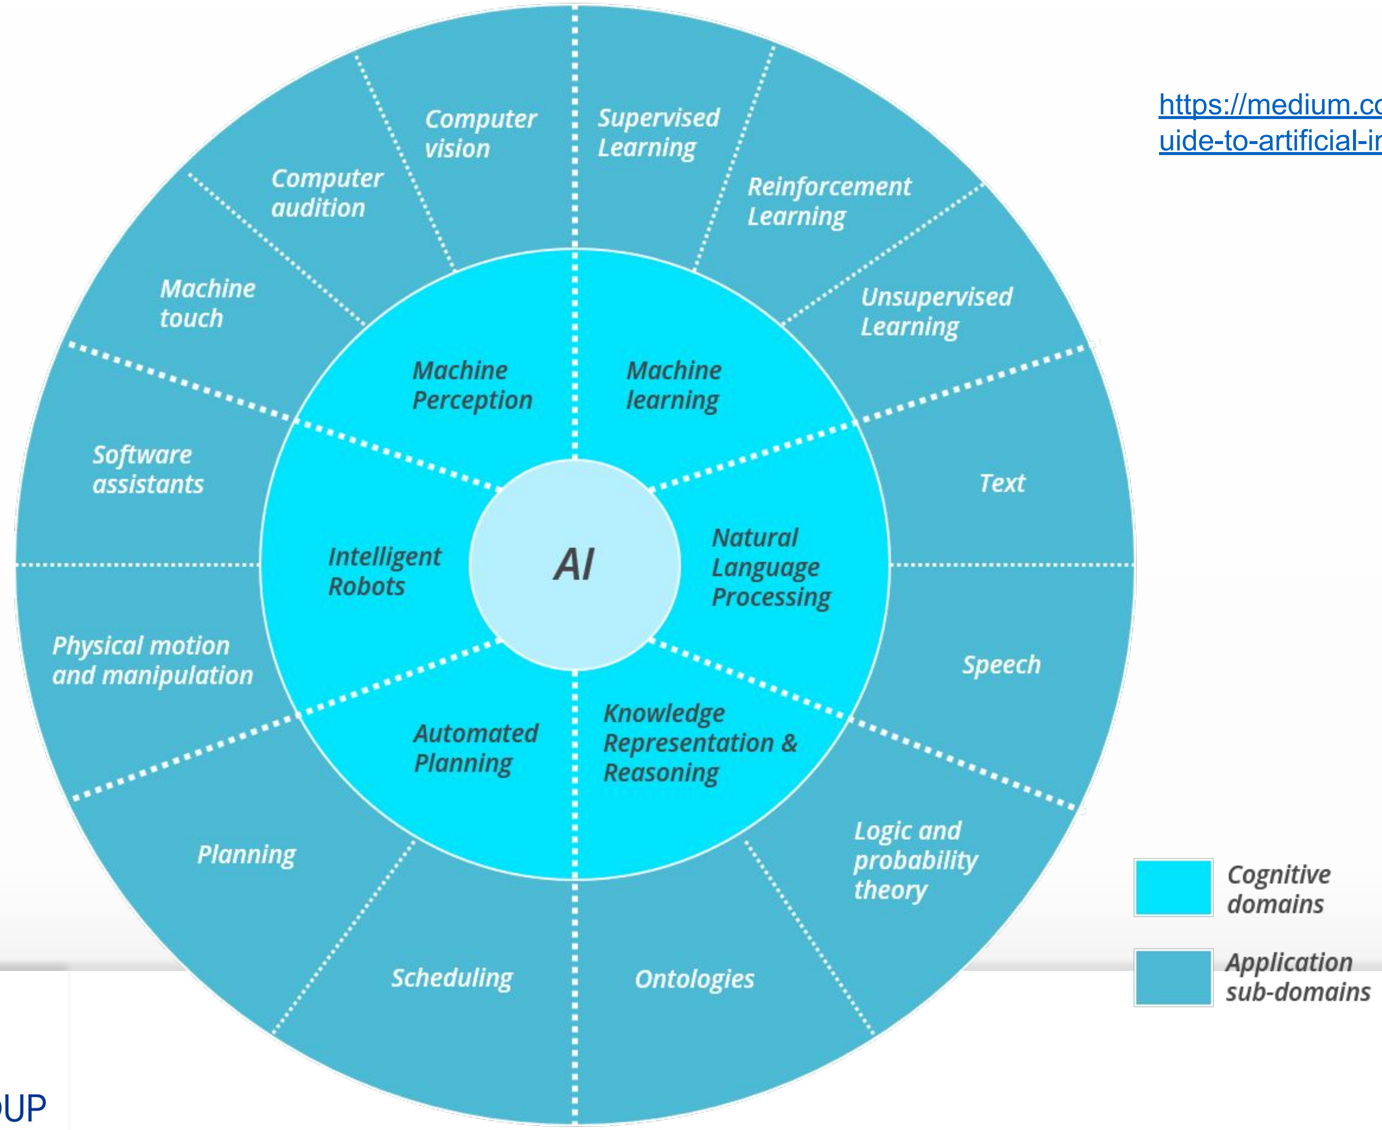

https://medium.com/appanion/a-five-minute-guide-to-artificial-intelligence-c4262be85fd3

Q#3
What is a chair?

Symbolic

AI is all about **representations**

Symbolic ("top-down"): good-old-fashioned AI

- **concepts** understandable by humans
- manipulation of these concepts
- good in creating **new knowledge** (in a well contained domain)

What is a **chair**? (part 2)

Subsymbolic (“bottom-up”)

- “raw” measurable **data**, e.g. from sensors
- **robust**, good for ill-defined (**simple**) tasks

Sub-symbolic **example**

Example: meeting room occupation

temperature IR sensor

We want to detect automatically when the classroom is occupied.

⇒ Set up an experiment to learn the relation!

First look at sensor data

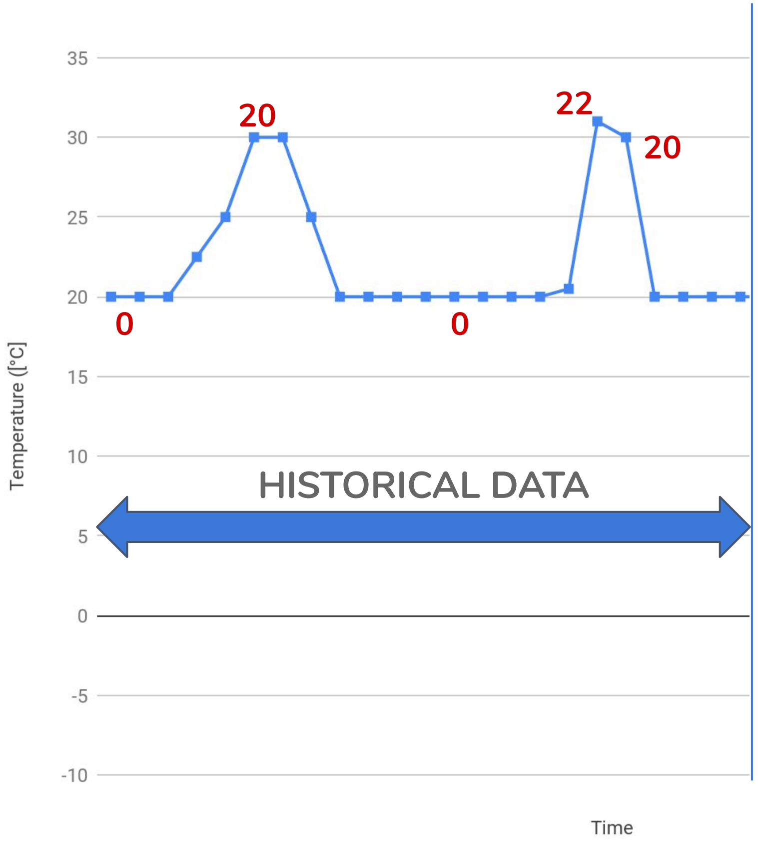

Nice: we can use AI to learn the #persons in the room.

Used for decision making:

Rooms with <20 users/day will be discarded.

First look at sensor data

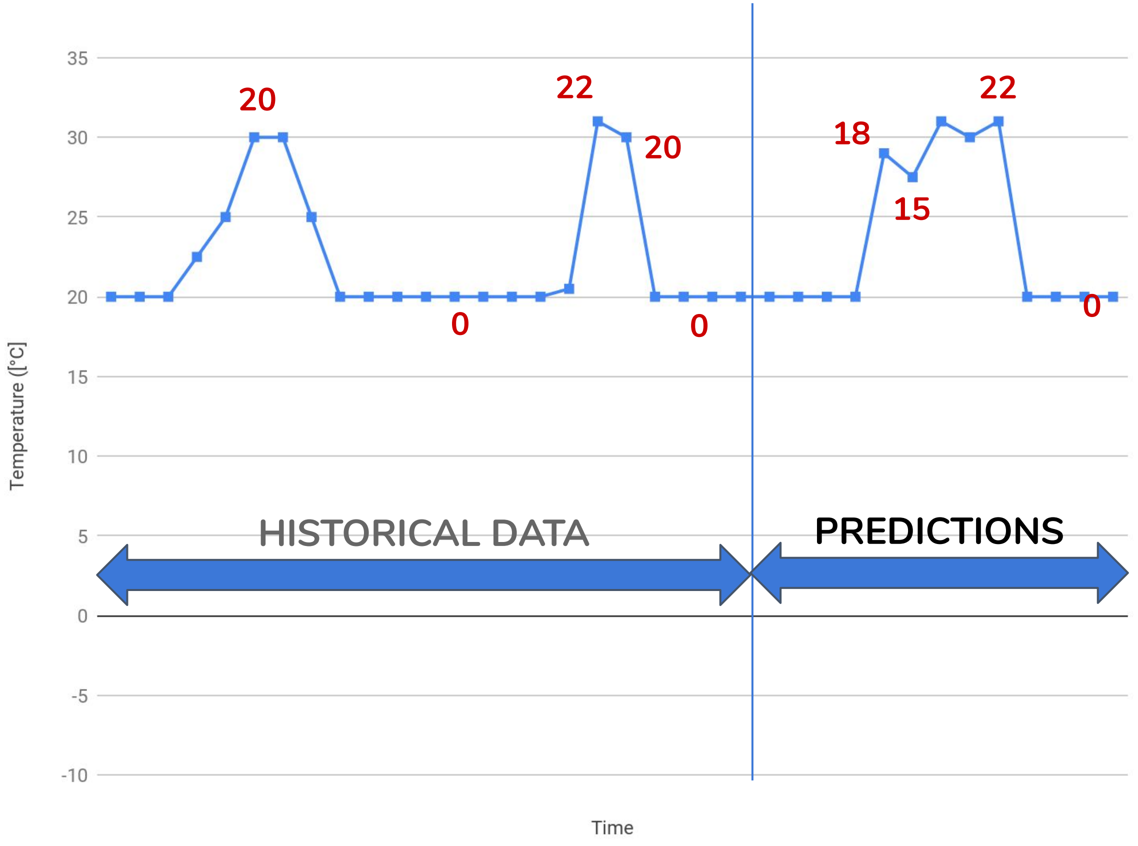

The board decides to **cut down energy.** What will happen to the predictions?

Energy saving ⇒ A/C off at night

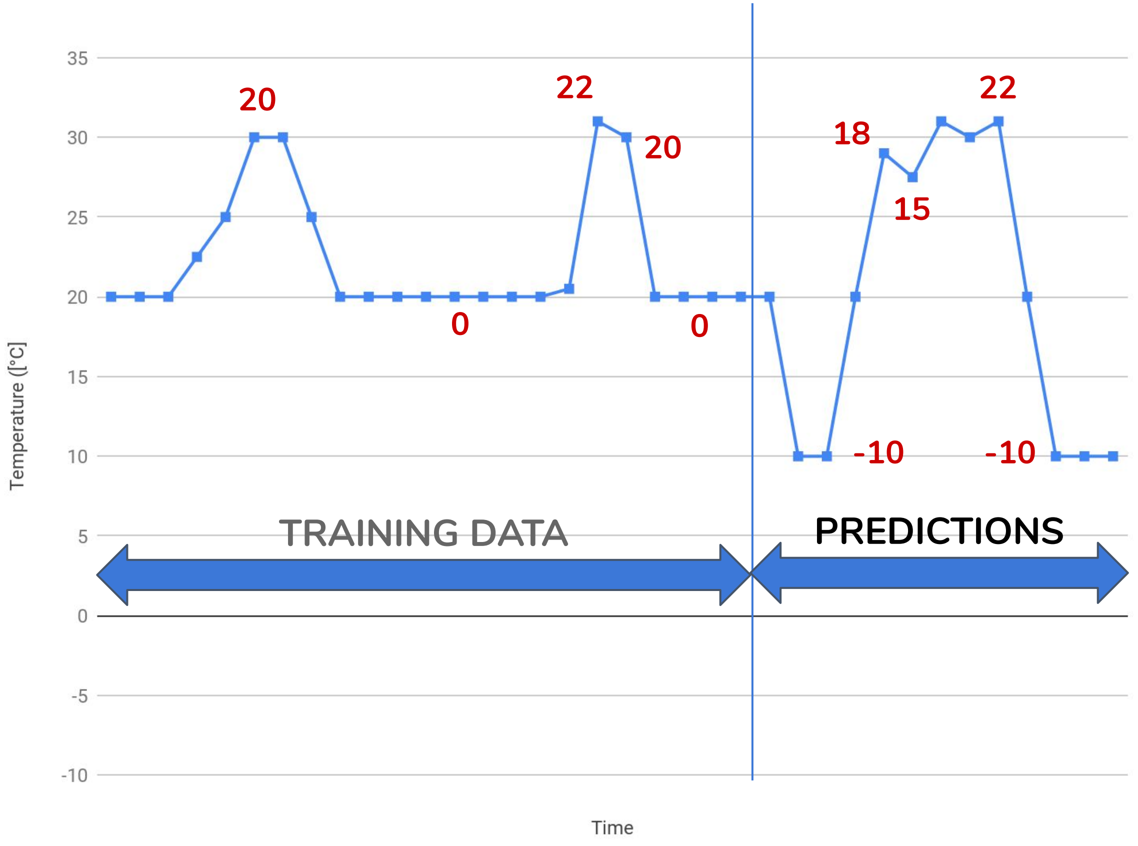

Negative #persons?! What is the consequence?

The model **overgeneralized.**

Room will be discarded as the sum < 20.

AHA!

We know that the #persons cannot be negative.

Strange... what happens here?

Room no longer in use?
Very cold outside?

The **sensor** input was **blocked** and outputs 0.

Room will be discarded.

Then summer comes...

A/C is shut off in summer.
20 °C is no longer a reference.

We know the sensor is just a proxy to measure t° increase by human presence

Algorithms have no idea of the context, unless you tell them so.

Data are a projection of reality

Which **representation** do you prefer to calculate a route?

Q#4

When did AI start?

What are the roots, foundations?

Plato

Reality

… how we **represent / think about** it

… and how we **perceive** it

Aristotle, Descartes, Leibniz

- We can represent reality as **symbols**
- And **reasoning** can represent our thinking
- **Machines** can perform reasoning

Ada Lovelace

[The Analytical Engine] might **act upon other things besides *number***, were objects found whose mutual fundamental relations could be expressed by those of the **abstract science of operations**...

Supposing, for instance, that the fundamental relations of pitched sounds in the science of harmony and of musical composition were susceptible of such expression and adaptations, **the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent.**

Russel, Gödel, Turing, Shannon

- Formal systems
- Turing machines
- Algorithms that work on representations to solve tasks
- Cybernetics, information theory, systems theory

(1950)

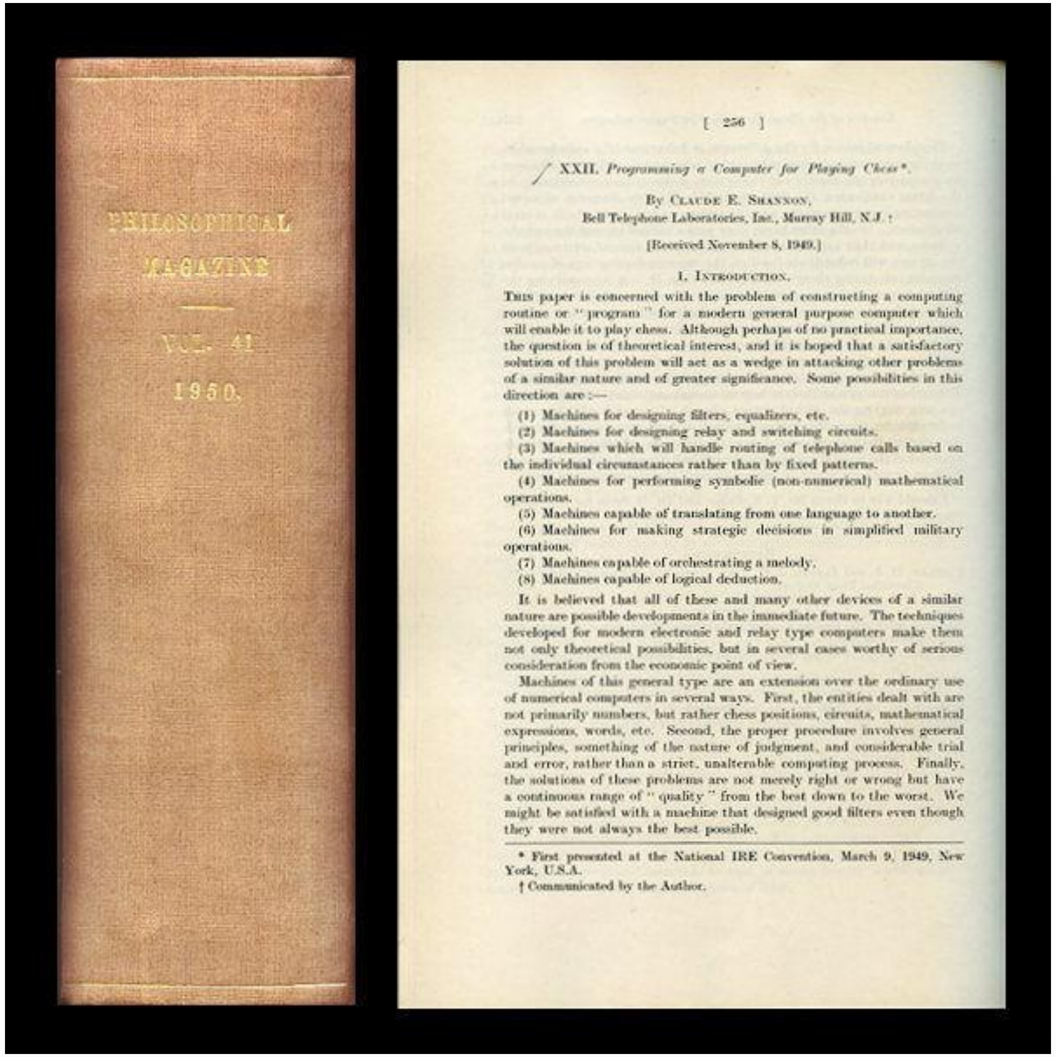

[256]

XXII. *Programming a Computer for Playing Chess**.

By CLAUDE E. SHANNON,
Bell Telephone Laboratories, Inc., Murray Hill, N.J.†

[Received November 8, 1949.]

I. INTRODUCTION.

THIS paper is concerned with the problem of constructing a computing routine or "program" for a modern general purpose computer which will enable it to play chess. Although perhaps of no practical importance, the question is of theoretical interest, and it is hoped that a satisfactory solution of this problem will act as a wedge in attacking other problems of a similar nature and of greater significance. Some possibilities in this direction are :—

(1) Machines for designing filters, equalizers, etc.
(2) Machines for designing relay and switching circuits.
(3) Machines which will handle routing of telephone calls based on the individual circumstances rather than by fixed patterns.
(4) Machines for performing symbolic (non-numerical) mathematical operations.
(5) Machines capable of translating from one language to another.
(6) Machines for making strategic decisions in simplified military operations.
(7) Machines capable of orchestrating a melody.
(8) Machines capable of logical deduction.

It is believed that all of these and many other devices of a similar nature are possible developments in the immediate future. The techniques developed for modern electronic and relay type computers make them not only theoretical possibilities, but in several cases worthy of serious consideration from the economic point of view.

Machines of this general type are an extension over the ordinary use of numerical computers in several ways. First, the entities dealt with are not primarily numbers, but rather chess positions, circuits, mathematical expressions, words, etc. Second, the proper procedure involves general principles, something of the nature of judgment, and considerable trial and error, rather than a strict, unalterable computing process. Finally, the solutions of these problems are not merely right or wrong but have a continuous range of "quality" from the best down to the worst. We might be satisfied with a machine that designed good filters even though they were not always the best possible.

* First presented at the National IRE Convention, March 9, 1949, New York, U.S.A.

† Communicated by the Author.

1. Zero-order approximation (symbols independent and equiprobable).

   XFOML RXKHRJFFJUJ ZLPWCFWKCYJ FFJEYVKCQSGHYD QPAAMKBZAACIBZL-HJQD.

2. First-order approximation (symbols independent but with frequencies of English text).

   OCRO HLI RGWR NMIELWIS EU LL NBNESEBYA TH EEI ALHENHTTPA OOBTTVA NAH BRL.

3. Second-order approximation (digram structure as in English).

   ON IE ANTSOUTINYS ARE T INCTORE ST BE S DEAMY ACHIN D ILONASIVE TU-COOWE AT TEASONARE FUSO TIZIN ANDY TOBE SEACE CTISBE.

4. Third-order approximation (trigram structure as in English).

   IN NO IST LAT WHEY CRATICT FROURE BIRS GROCID PONDENOME OF DEMONS-TURES OF THE REPTAGIN IS REGOACTIONA OF CRE.

5. First-order word approximation. Rather than continue with tetragram, . . . , *n*-gram structure it is easier and better to jump at this point to word units. Here words are chosen independently but with their appropriate frequencies.

   REPRESENTING AND SPEEDILY IS AN GOOD APT OR COME CAN DIFFERENT NAT-URAL HERE HE THE A IN CAME THE TO OF TO EXPERT GRAY COME TO FURNISHES THE LINE MESSAGE HAD BE THESE.

6. Second-order word approximation. The word transition probabilities are correct but no further structure is included.

   THE HEAD AND IN FRONTAL ATTACK ON AN ENGLISH WRITER THAT THE CHAR-ACTER OF THIS POINT IS THEREFORE ANOTHER METHOD FOR THE LETTERS THAT THE TIME OF WHO EVER TOLD THE PROBLEM FOR AN UNEXPECTED.

(1948)

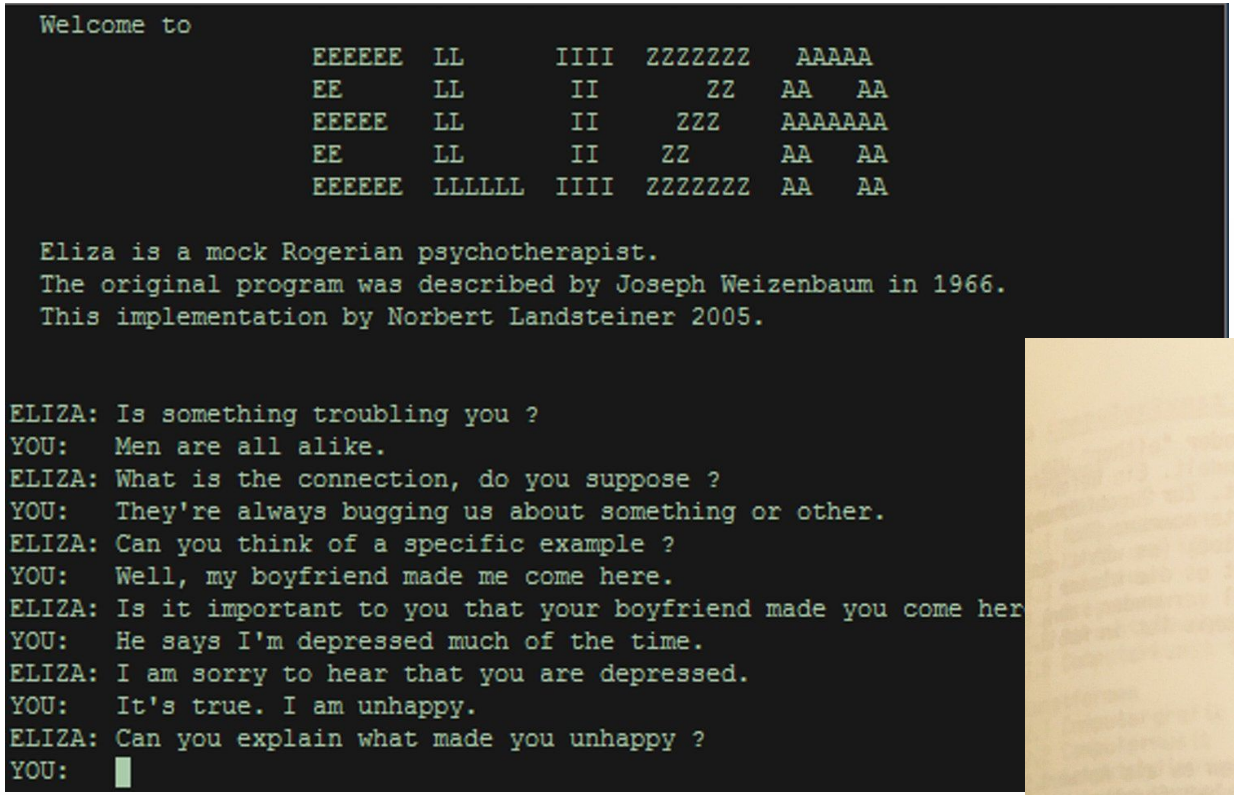

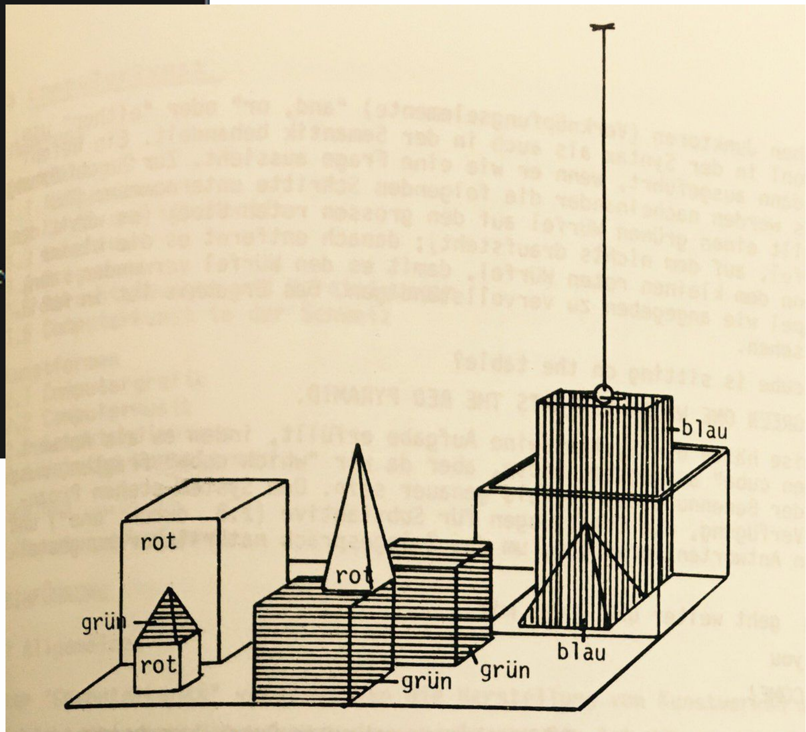

Abbildung 3: "Find a block which is taller than the one you are holding and put it into the box".

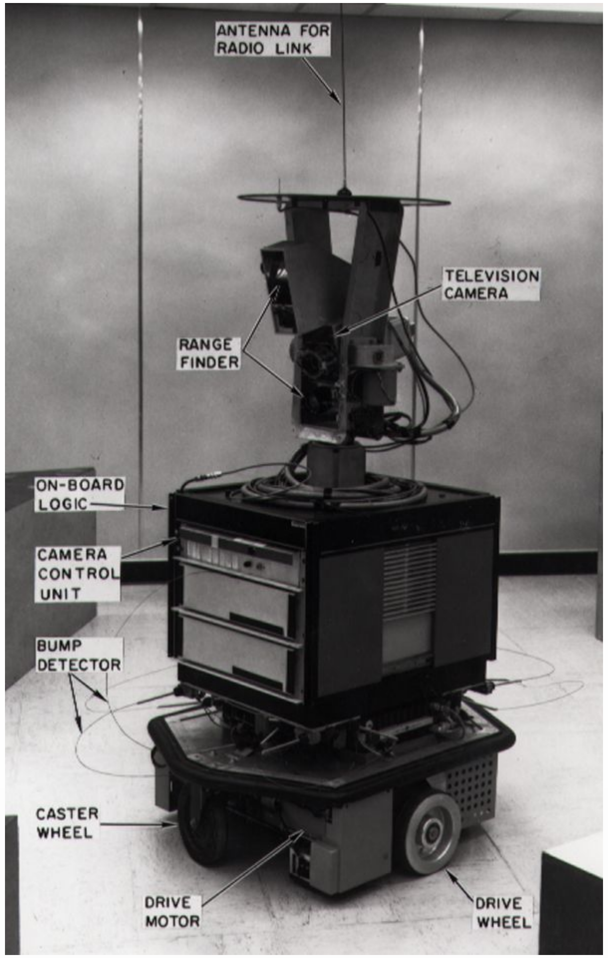

Q#5
What makes AI hard?

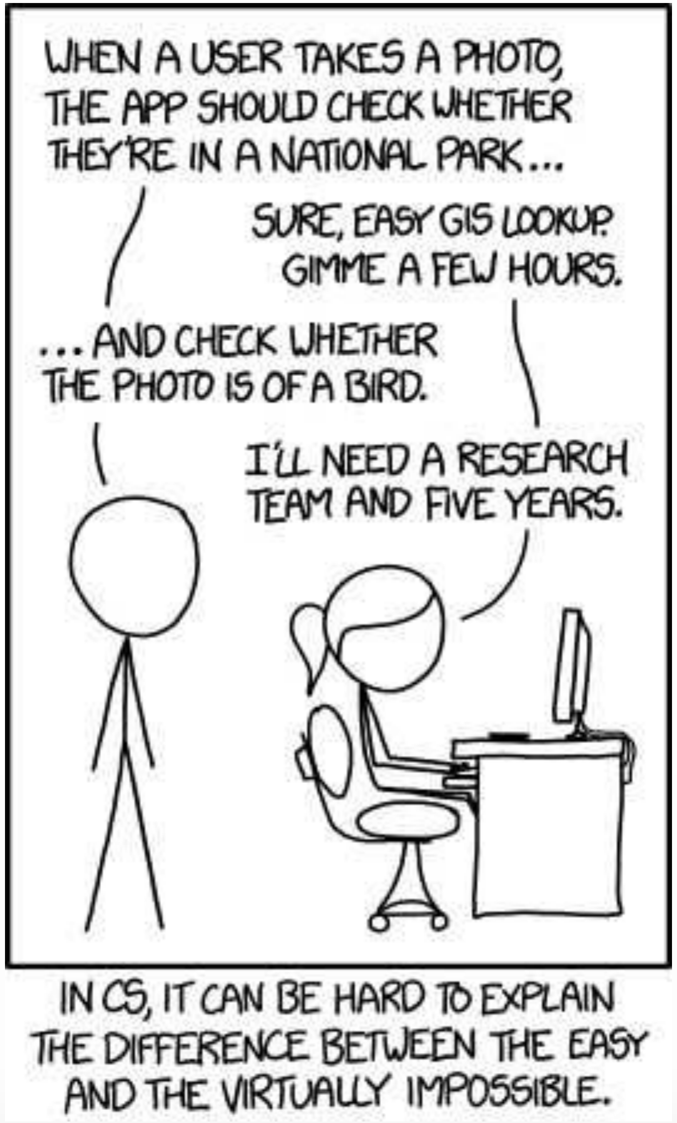
WHEN A USER TAKES A PHOTO, THE APP SHOULD CHECK WHETHER THEY'RE IN A NATIONAL PARK...
SURE, EASY GIS LOOKUP. GIMME A FEW HOURS.
...AND CHECK WHETHER THE PHOTO IS OF A BIRD.
I'LL NEED A RESEARCH TEAM AND FIVE YEARS.
IN CS, IT CAN BE HARD TO EXPLAIN THE DIFFERENCE BETWEEN THE EASY AND THE VIRTUALLY IMPOSSIBLE.

Interdisciplinary (intersection, not union)

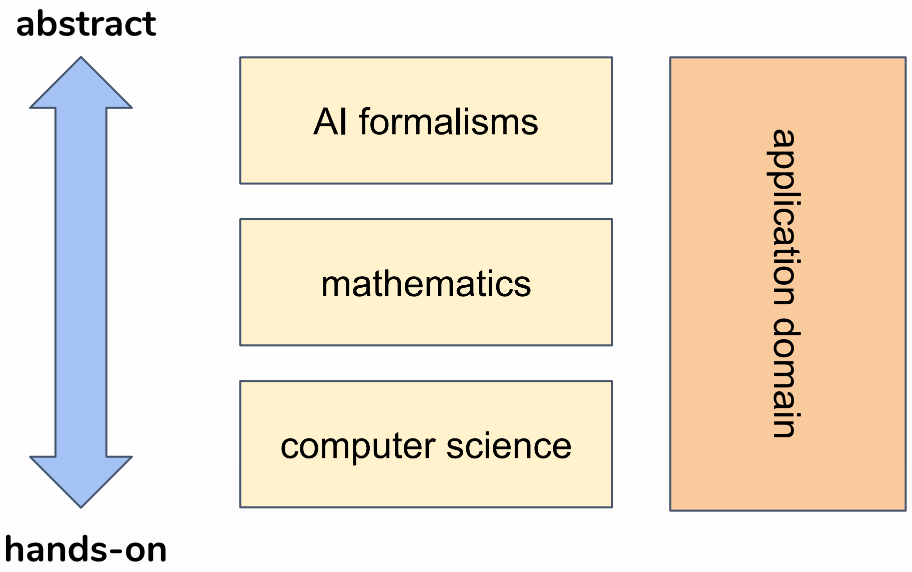

how does my business work?

but explained **to** a computer…

that is, **in mathematical terms?**

“Understanding” how it works

- conceptually?
- mathematically?
- algorithmically?

- Empirical risk minimization
  - framework to design learning algorithms
  - $l(f(\mathbf{x}^{(t)};$
  - $\Omega(\boldsymbol{\theta})$ is a
- Learning i
  - ideally, we'
  - loss function is a surrogate for what we truly should optimize (e.g. upper bound)

```
# Hyperparameters for our network
input_size = 784
hidden_sizes = [128, 64]
output_size = 10

# Build a feed-forward network
model = nn.Sequential(nn.Linear(input_size, hidden_sizes[0]),
                      nn.ReLU(),
                      nn.Linear(hidden_sizes[0], hidden_sizes[1]),
                      nn.ReLU(),
                      nn.Linear(hidden_sizes[1], output_size),
                      nn.Softmax(dim=1))
print(model)
```

space

It challenges our beliefs.

#1 - A walking robot should be stable at all times.

AI continuously **challenges** our **beliefs** and **knowledge**, as it learns to reproduce our results, and will take over a big part of **experimentation** too.

#2 - Human & algorithmic biases

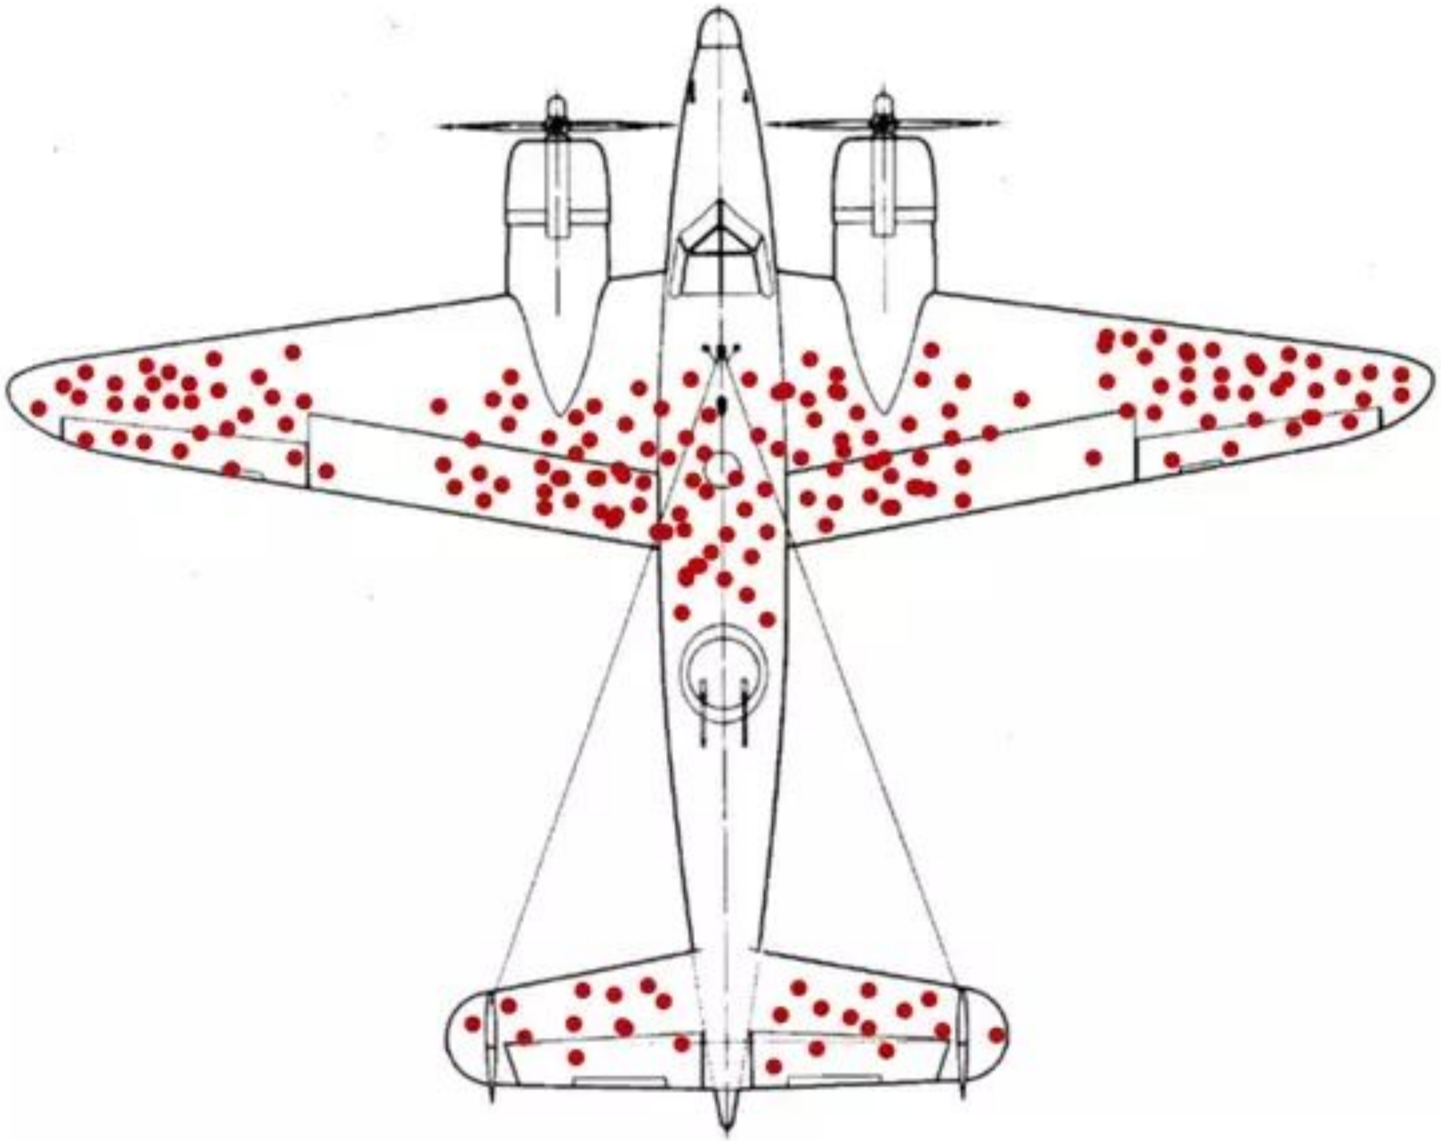

Bullets in airplanes that came back in WW II.

Where to reinforce the armour?

How many f’s in the following sentence?

“Finished files are the result of years of scientific study combined with the experience of years.”

#3 - What action should I take?

Language & reality is **ambiguous**

(a chair, you said?)

“Descendez-le!”

Context matters

- “Could you please pass me the glass of water?”

when I’m thirsty?

when I’m holding a bottle of water?

ARTIFICIAL INTELLIGENCE RESEARCH GROUP

#4 - Where is the flag?

We **reason** over language

the flag is **left** of the table!

the flag is **right** of the table!

“Hybrid” AI: data + reasoning

Q#6
Why do we have a revival now?

Why a **revival?**

- First **digitalization** wave has happened
  (**data** can be harvested, digital **interfaces** to processes & users, SOA)
- **Global** interconnected & **servitized** world & market
  (more **competition**, higher user expectations, new business models)
- **Cloud** services, **open source** & computing **power**
  (powerful, modular **building blocks**, small can compete with big)
- **Consumer driven** mindset
  (users compare with the best of *all,* not best of category)

Technology is moving to end-users

- Human interfaces to services are disappearing
  - bank tellers ⇒ ATMs
  - bank clerks ⇒ apps

- Users thus interact directly with **technology**.
  - Still very unnatural compared to humans (e.g. MOOC)
  - Quality and comfort are crucial
  - Digital natives

Q#7
What will AI change?

What is the **impact** of fixing three numbers?

Answer: **global prosperity**

research
pessimists

sale
optimists

AI is transformative

AI is **transformative**

1. **Changes the *nature* & *discovery* of knowledge**
   a. Allows companies to formalize knowledge
   b. Moves from humans to self-learning machines
   c. Driving force in science, the new “statistics”

2. Enables **autonomous** and **adaptive** systems
   a. humans go from commanding to **interacting,** giving up control
   b. go beyond our understanding (e.g. high dimensional / nonlinear)

3. They have “infinite” **scale**.
   a. micro-personalization, human empowerment
   b. winner-takes-all economies as you can copy the best

We have come to a point that the systems we need
are **too complex for humans** to build, understand or maintain.

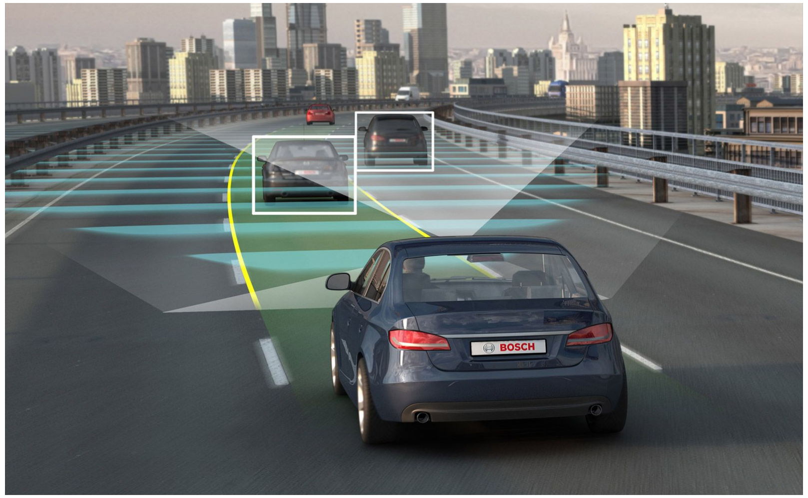

The real transformative value of AI lies exactly in the systems we do not understand

Explainability

Transparency of algorithms?

$$f_c(z) = z^2 + c$$

AlphaGo

1. **Better** at Go than humans
   (scientists thought this breakthrough to happen in 2030 earliest)

2. Learn to play **in 2 days**

3. And “invent” **creative moves**

4. An **infinite** #players can be created instantaneously

AI advances science

- **Automate** tedious research tasks
- **Detect** complex patterns, rules… “New statistics”
- Improve **Randomized Controlled Trials**
- The **user is the experiment** => longitudinal studies, we learn a lot about humans

Unpredictable impact

Humans vs. machines

What happened to…

- bank tellers
- video stores
- accountants
- webmasters

And so who are…

- content managers
- community managers
- IT car repair men
- elevator service personnel

What will happen to?

- travel agents
- real estate agents
- broadcasters
- farmers
- marketeers

Have we become the robots?

- algorithms orchestrating people in warehouses;
- opening doors to flats for people decided by airbnb;
- buying books Amazon tells us to.

"General consensus"

- Will **change our jobs**
  - probably not eliminate, but change the tasks we do
  - and create new jobs!
- Our **interaction with products** will change
  - systems that have their own plan, "will"
- **Lifelong learning** becoming main learning mode
- **Polarization** of tech, data & talent

Q#8
Why is it freaky?

buy these publications, mail orders for these, give this "public opinion"

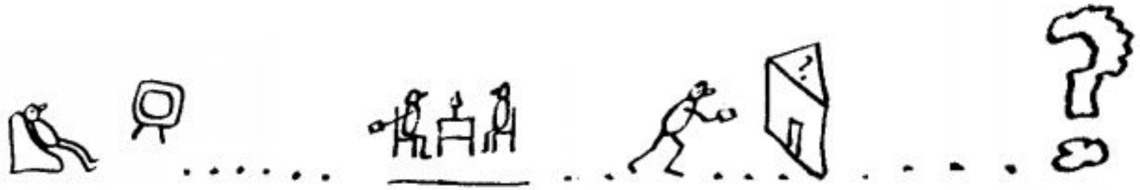

watch these shows, use this credit card? and this? to buy this

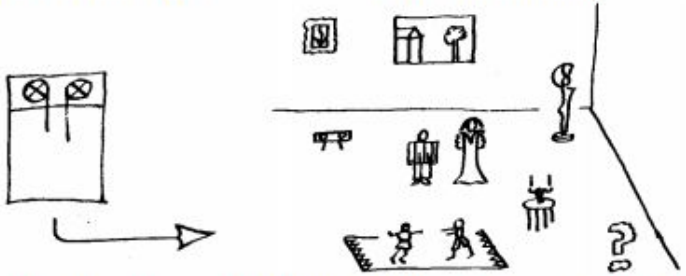

and is therefore exactly like this with. . . certain susceptiblities. . .

Good Morning! Your new life style! New politics! New religion! How can I ever thank you?

Creative ways…

Turn your competitors' visitors into your customers

By combining location intelligence with online advertising potential OMNIcookie empowers retailers to reconnect with shoppers after they've left the store

SEE HOW IT WORKS

+60%
offline visits uplift

Biggest brands
working with the best

+1 billion
devices in Asia alone

No installation
no hardware, no beacons

Vacuum cleaners

What does your interior design tell about

- your personality?
- your lifestyle?

Also when “logged out” (passive authentication)

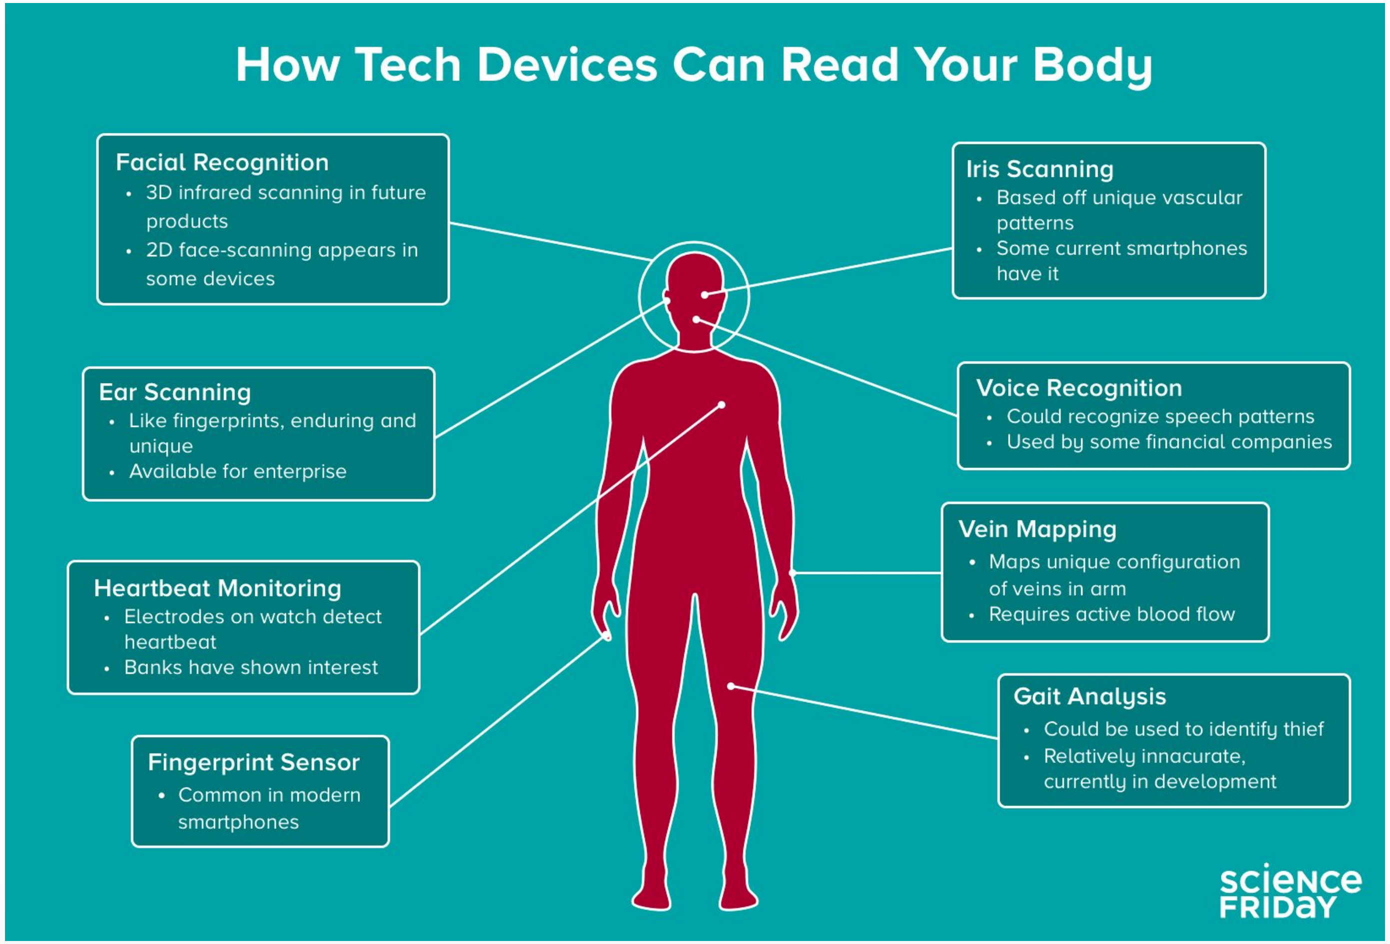

Personal identification

People can now be identified at a distance by their heartbeat

And then dealt with, if they are enemy operatives

ARTIFICIAL INTELLIGENCE RESEARCH GROUP

https://witanworld.com/article/2018/08/12/witan-sapience-biometrics/
https://www.economist.com/science-and-technology/2020/01/23/people-can-now-be-identified-at-a-distance-by-their-heartbeat
https://seon.io/sense-platform/

@aibrussels

Recommender systems

Winner technology of AI

- What advertisements you see (Google)
- What movie you watch (Netflix)
- What products you buy (Amazon)
- What news you read (Google, Facebook,…)
- What friends you make(Facebook, Linkedin,…)
- Who will be your spouse (Match.com, …)

The most pervasive (and profitable) of all AI applications.

Psychometrics

- Understanding the customer
- Classification into **specific** categories
- Prediction of specific quantities
- Clustering of your profile (similarity)

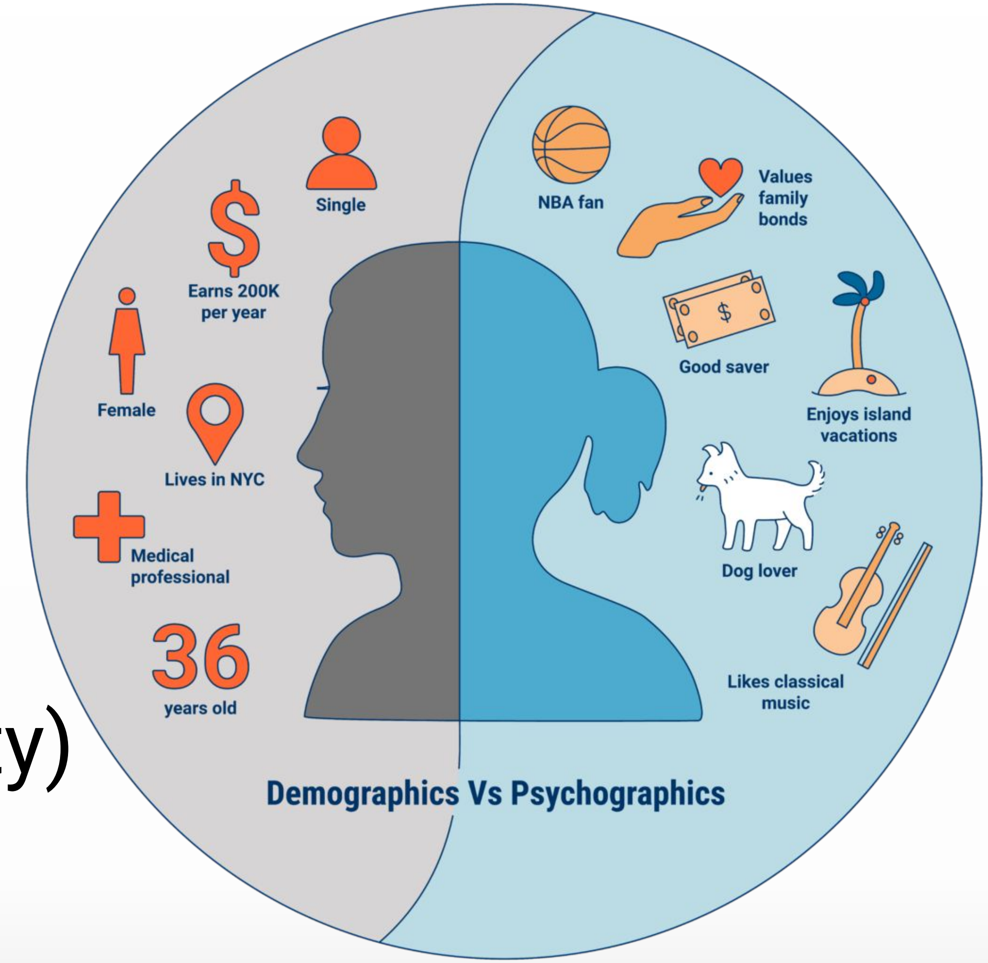

Source: CB insights

Saying someone's behaviour is extravert

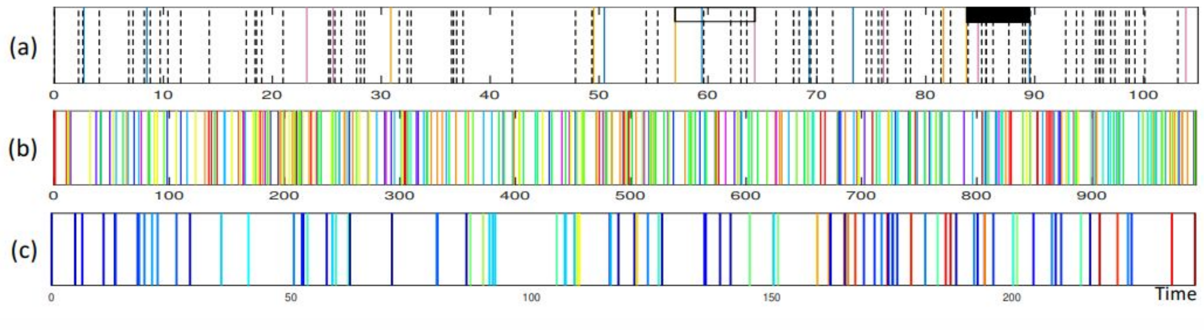

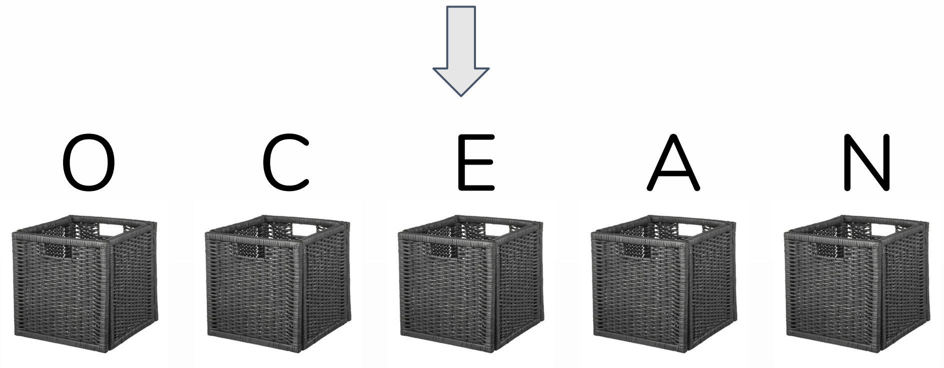

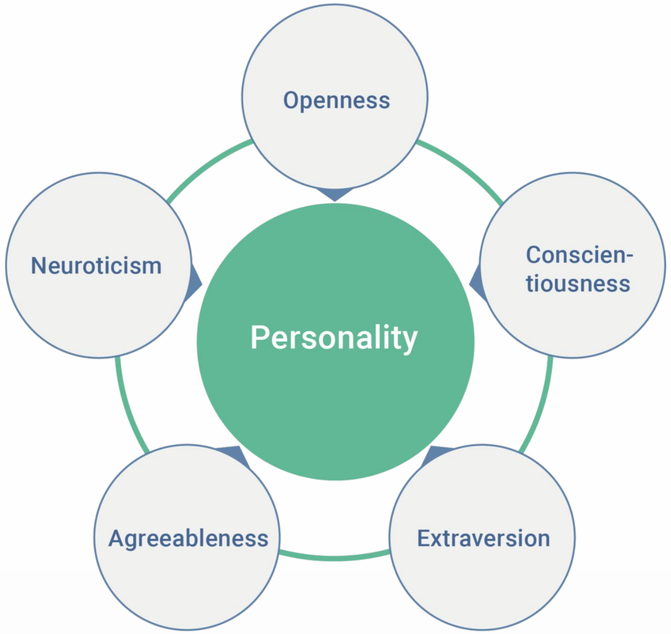

… is different from understanding what that means, or **reason** about it.

Examples

Uber: Users Are More Likely To Pay Surge Pricing If Their Phone Battery Is Low

Amit Chowdhry Contributor

Tech enthusiast, born in Ann Arbor and educated at Michigan State

Mac users pay more than PC users, says Orbitz

The travel site says Mac users will pay $20 to $30 a night more on hotels than PC users.

source: https://www.forbes.com/sites/amitchowdhry/2016/05/25/uber-low-battery/#3935789574b3

Polarization…

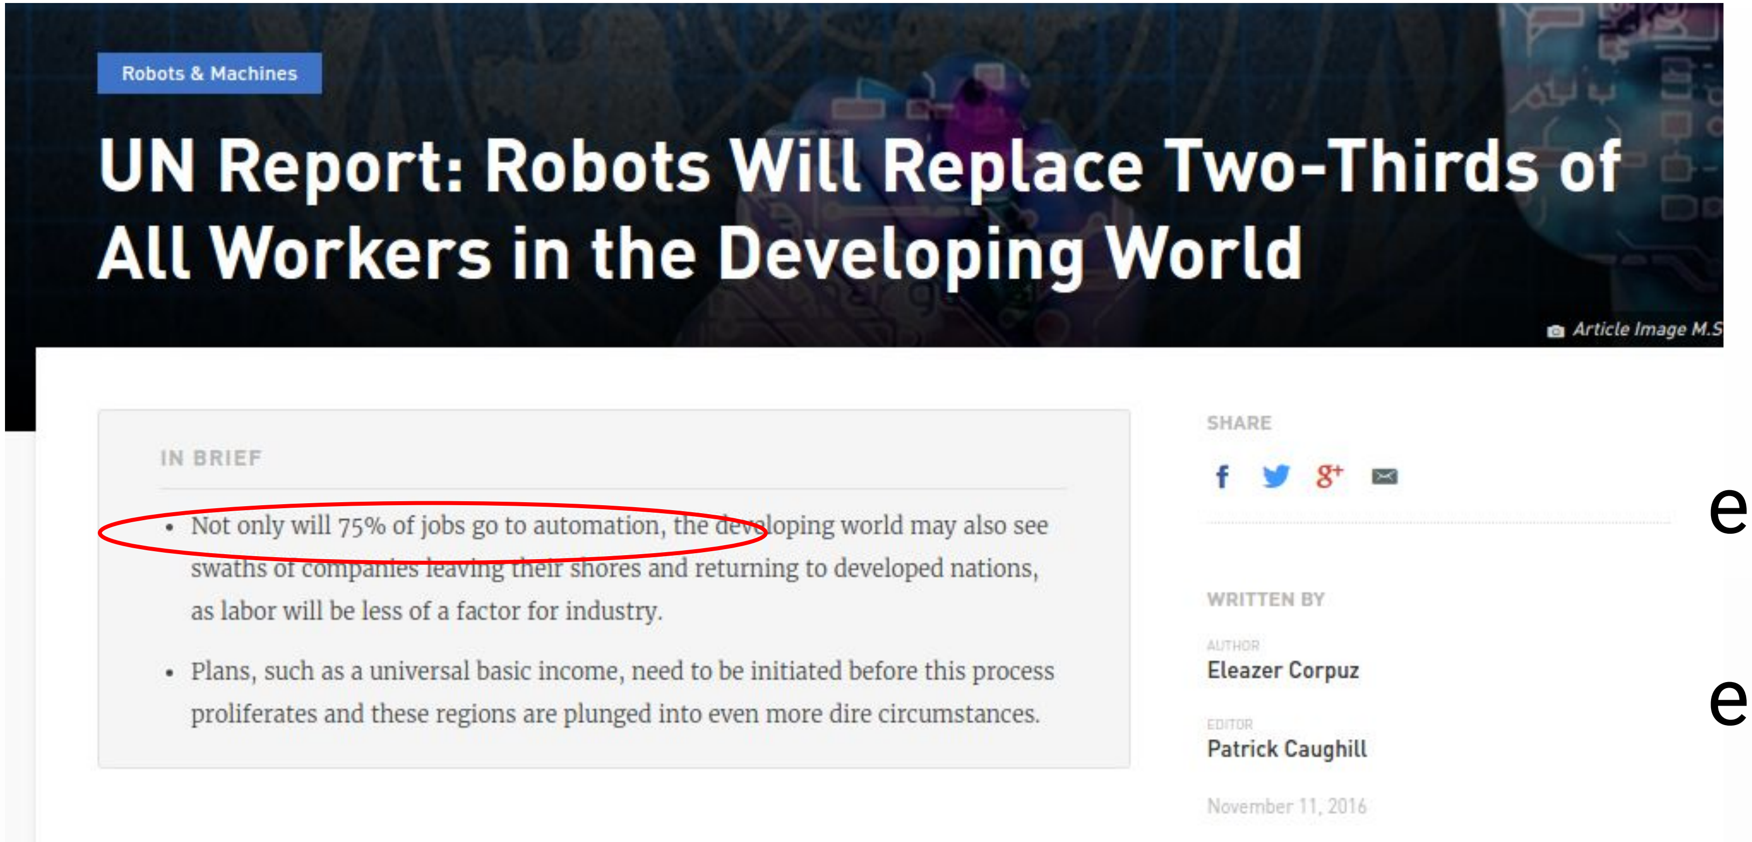

Robots & Machines

UN Report: Robots Will Replace Two-Thirds of All Workers in the Developing World

Article Image M.S

IN BRIEF

- Not only will 75% of jobs go to automation, the developing world may also see swaths of companies leaving their shores and returning to developed nations, as labor will be less of a factor for industry.
- Plans, such as a universal basic income, need to be initiated before this process proliferates and these regions are plunged into even more dire circumstances.

SHARE

WRITTEN BY

AUTHOR
Eleazer Corpuz

EDITOR
Patrick Caughill

November 11, 2016

education?

equality?

Filter bubble

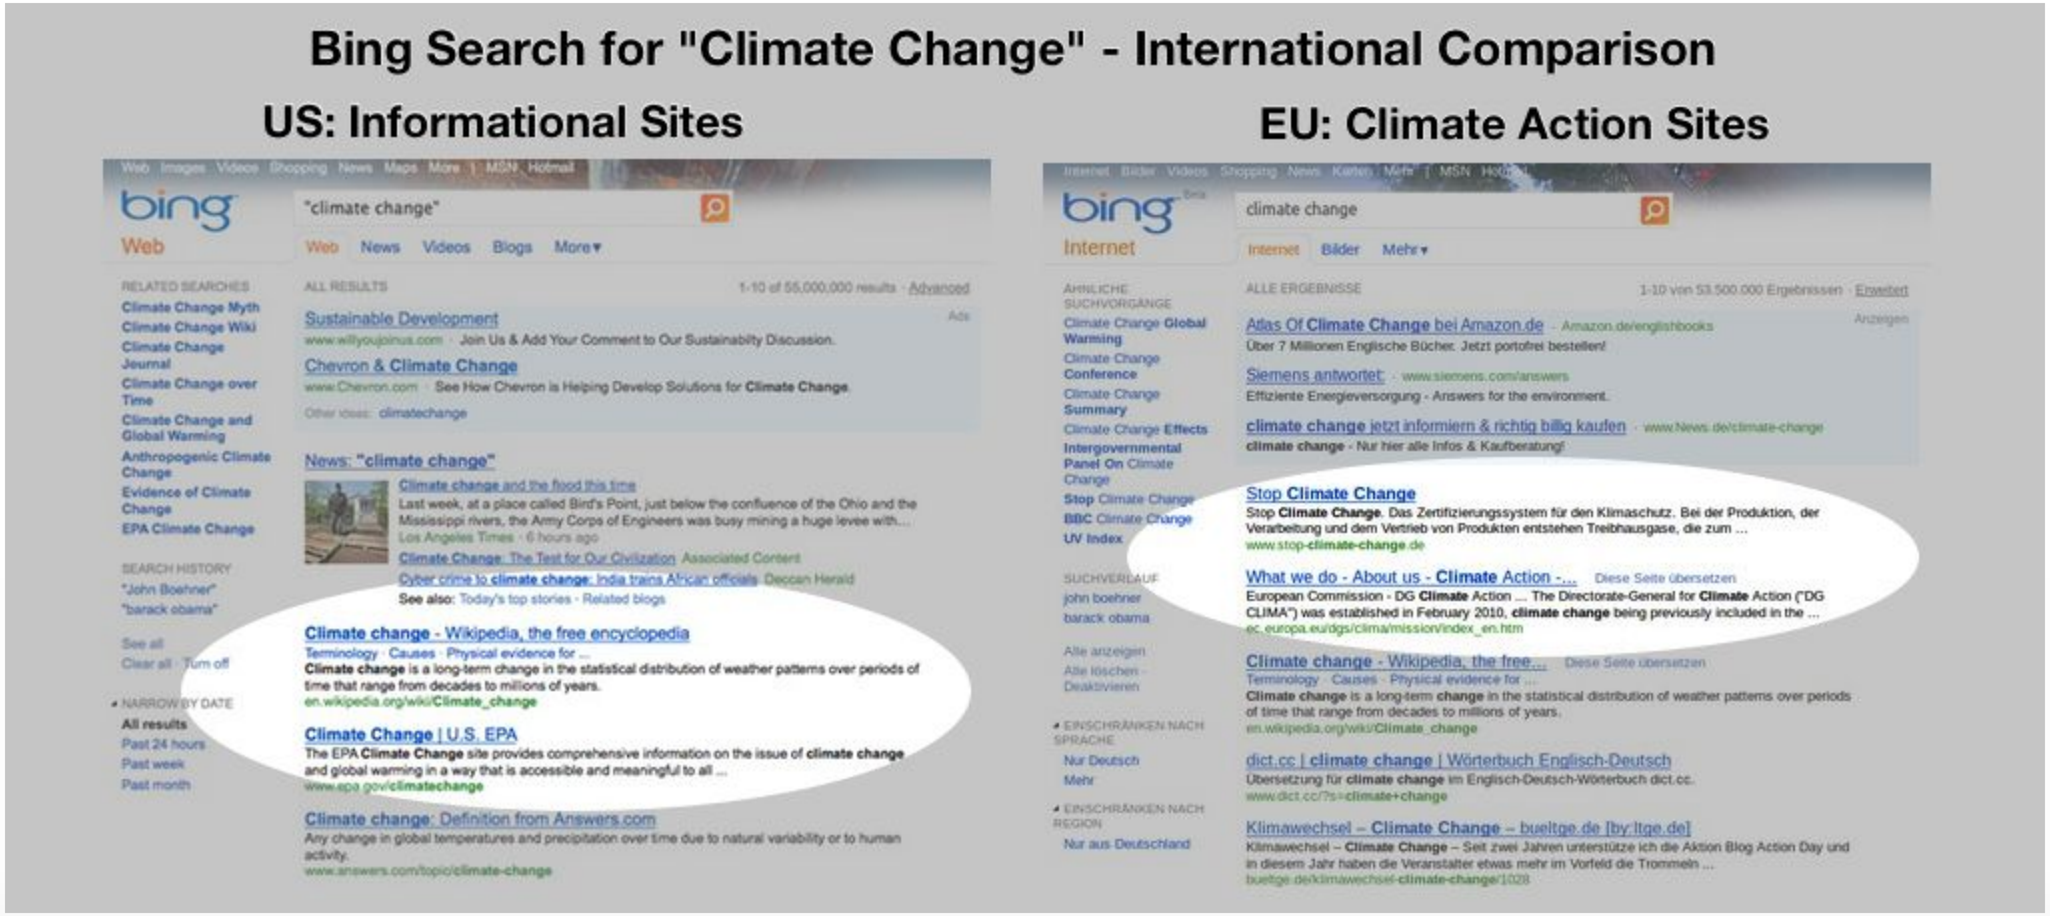

ARTIFICIAL INTELLIGENCE RESEARCH GROUP

source: https://www.slideshare.net/KiranGarimella1/quantifying-and-bursting-the-online-filter-bubble

@aibrussels

Polarization

- … of opinions
- … of people you meet
- … music you listen to
- … clothes you wear?

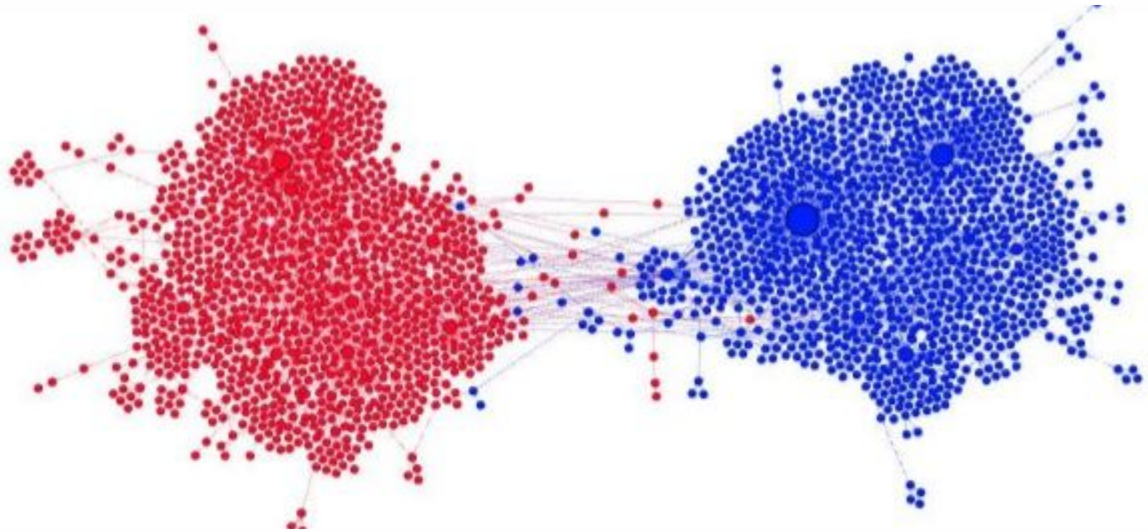

Deepfakes

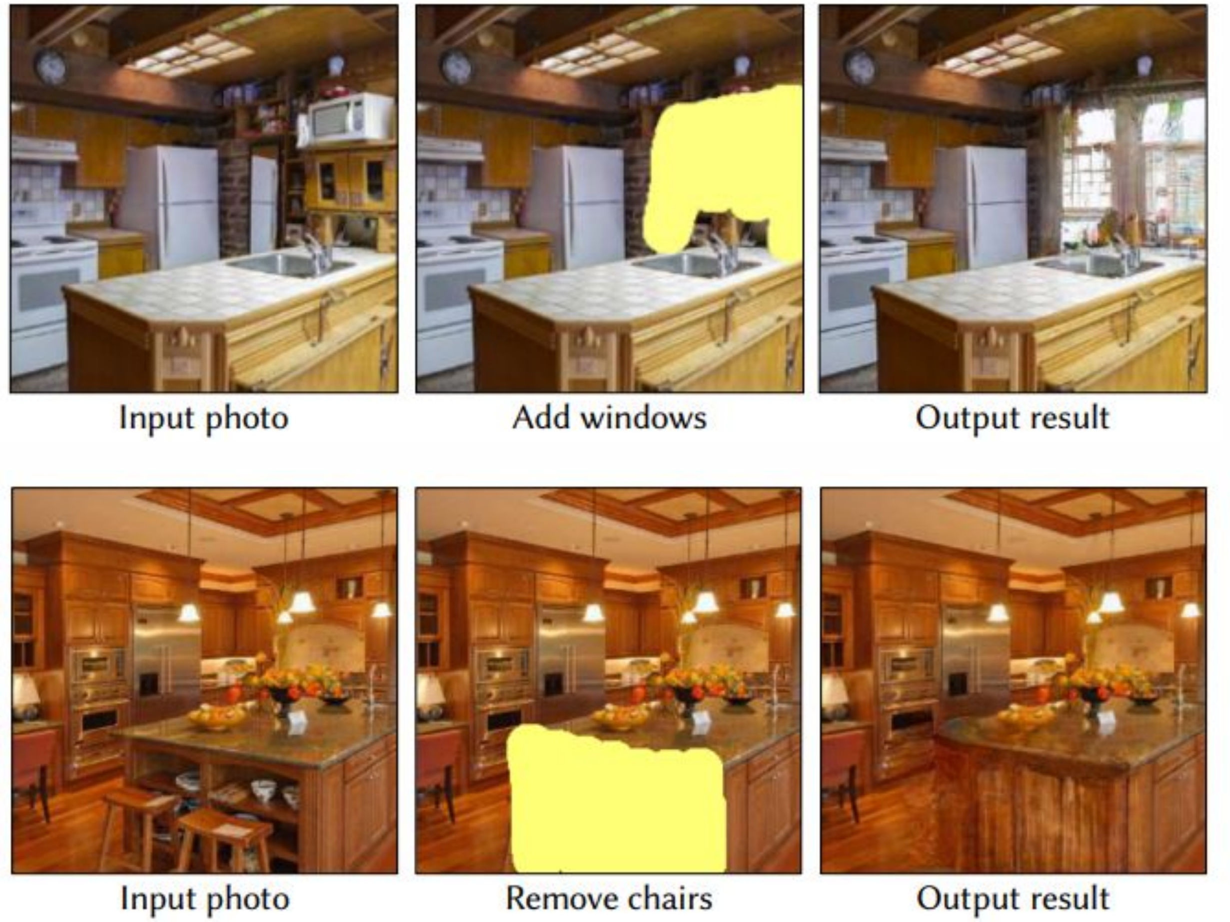

https://machinelearningmastery.com/impressive-applications-of-generative-adversarial-networks/

Example of Sketches to Color Photographs With pix2pix.Taken from Image-to-Image Translation with Conditional Adversarial Networks, 2016.

xample of GAN-Generated Photographs of Human PosesTaken from Pose Guided Person Image Generation, 2017.

Q#9
What can AI NOT do today?

"We are only 20 years away from a world in which machines will do any work a man can do"

H. Simon (1965)

Dartmouth summer program

- The term “Artificial Intelligence” was coined.
  - John McCarthy
  - Marvin Minsky
  - Nathaniel Rochester
  - Claude E. Shannon

- They expected to need **2 months** and a team of 10 to solve the problem

#1
we underestimate our own intelligence

We **underestimate** our intelligence!

- Tactile feedback
- Substance identification
- Memory of weight
- Geometric reasoning

Recognizing a guitar is different from

PHYSICAL

understanding that you can play music on it, or that if you can play on a guitar, you can play on a banjo.

#2
algorithms are ruthless

Algorithms are ruthless

Cobra effect: government policy to get rid of cobras. If a citizen brings in a venomous snake, they got a reward. People started breeding snakes.

Algorithms are ruthless (and so they **cheat**)

Algorithms Acting Out

- **Infanticide:** In a survival simulation, one AI species evolved to subsist on a diet of its own children.
- **Space War:** Algorithms exploited flaws in the rules of the galactic videogame *Elite Dangerous* to invent powerful new weapons.
- **Body Hacking:** A four-legged virtual robot was challenged to walk smoothly by balancing a ball on its back. Instead, it trapped the ball in a leg joint, then lurched along as before.
- **Goldilocks Electronics:** Software evolved circuits to interpret electrical signals, but the design only worked at the temperature of the lab where the study took place.
- **Optical Illusion:** Humans teaching a gripper to grasp a ball accidentally trained it to exploit the camera angle so that it appeared successful—even when not touching the ball.

CHEATING AI CAUGHT HIDING DATA USING STEGANOGRAPHY

by: Inderpreet Singh

63 Comments

January 3, 2019

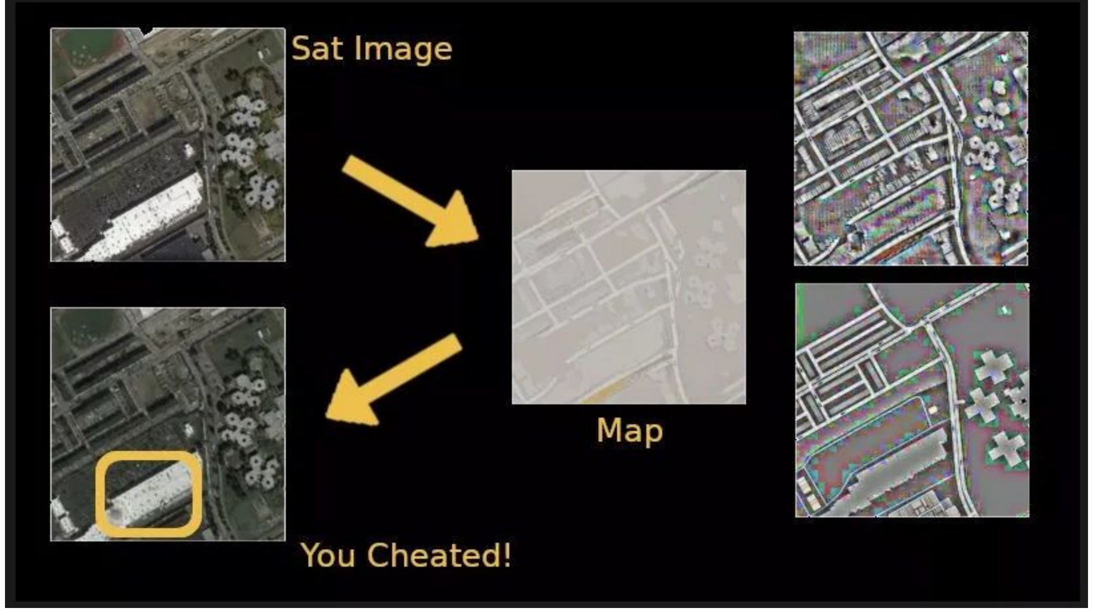

#3
humans are the limit:
we have to describe our problem mathematically

Formulating the problem = hard part

- It is like **explaining** how to
  - drive a car
  - cook
  - diagnose a patient
  - play chess
- **to a computer,** that has no understanding of context
- **MATHEMATICALLY!**

The **description** is the limit

The **complexity** of

- rules
- annotation
- reward signal

to describe the **problem**

that a **human** can handle is the sky.

Q#10
How would you
design a system that...

Answers questions like

Are there more blue things than yellow ones?

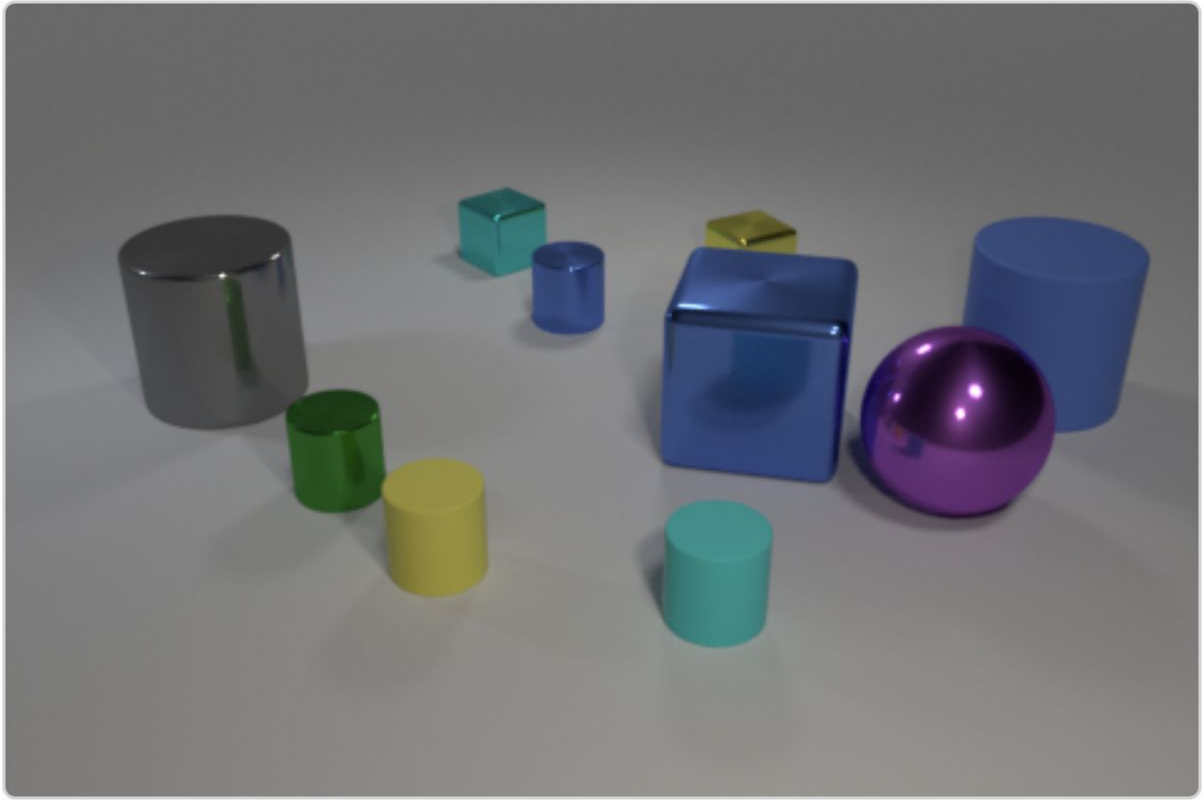

Visual question answering

https://ehai.ai.vub.ac.be/demos/visual-question-answering/

https://ehai.ai.vub.ac.be/demos/clevr-grammar/

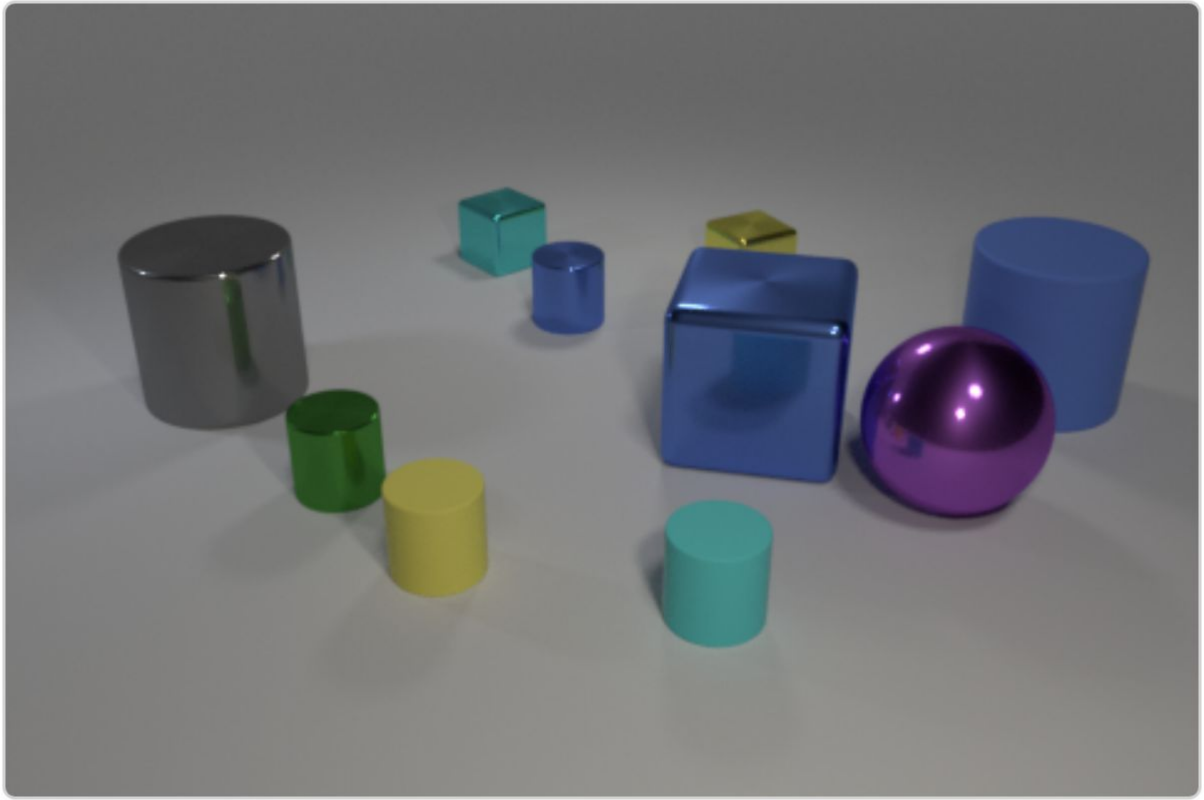

Questions?

Johan Loeckx
jloeckx@ai.vub.ac.be
https://ai.vub.ac.be

+32(0)486 37 98 71
@aibrussels