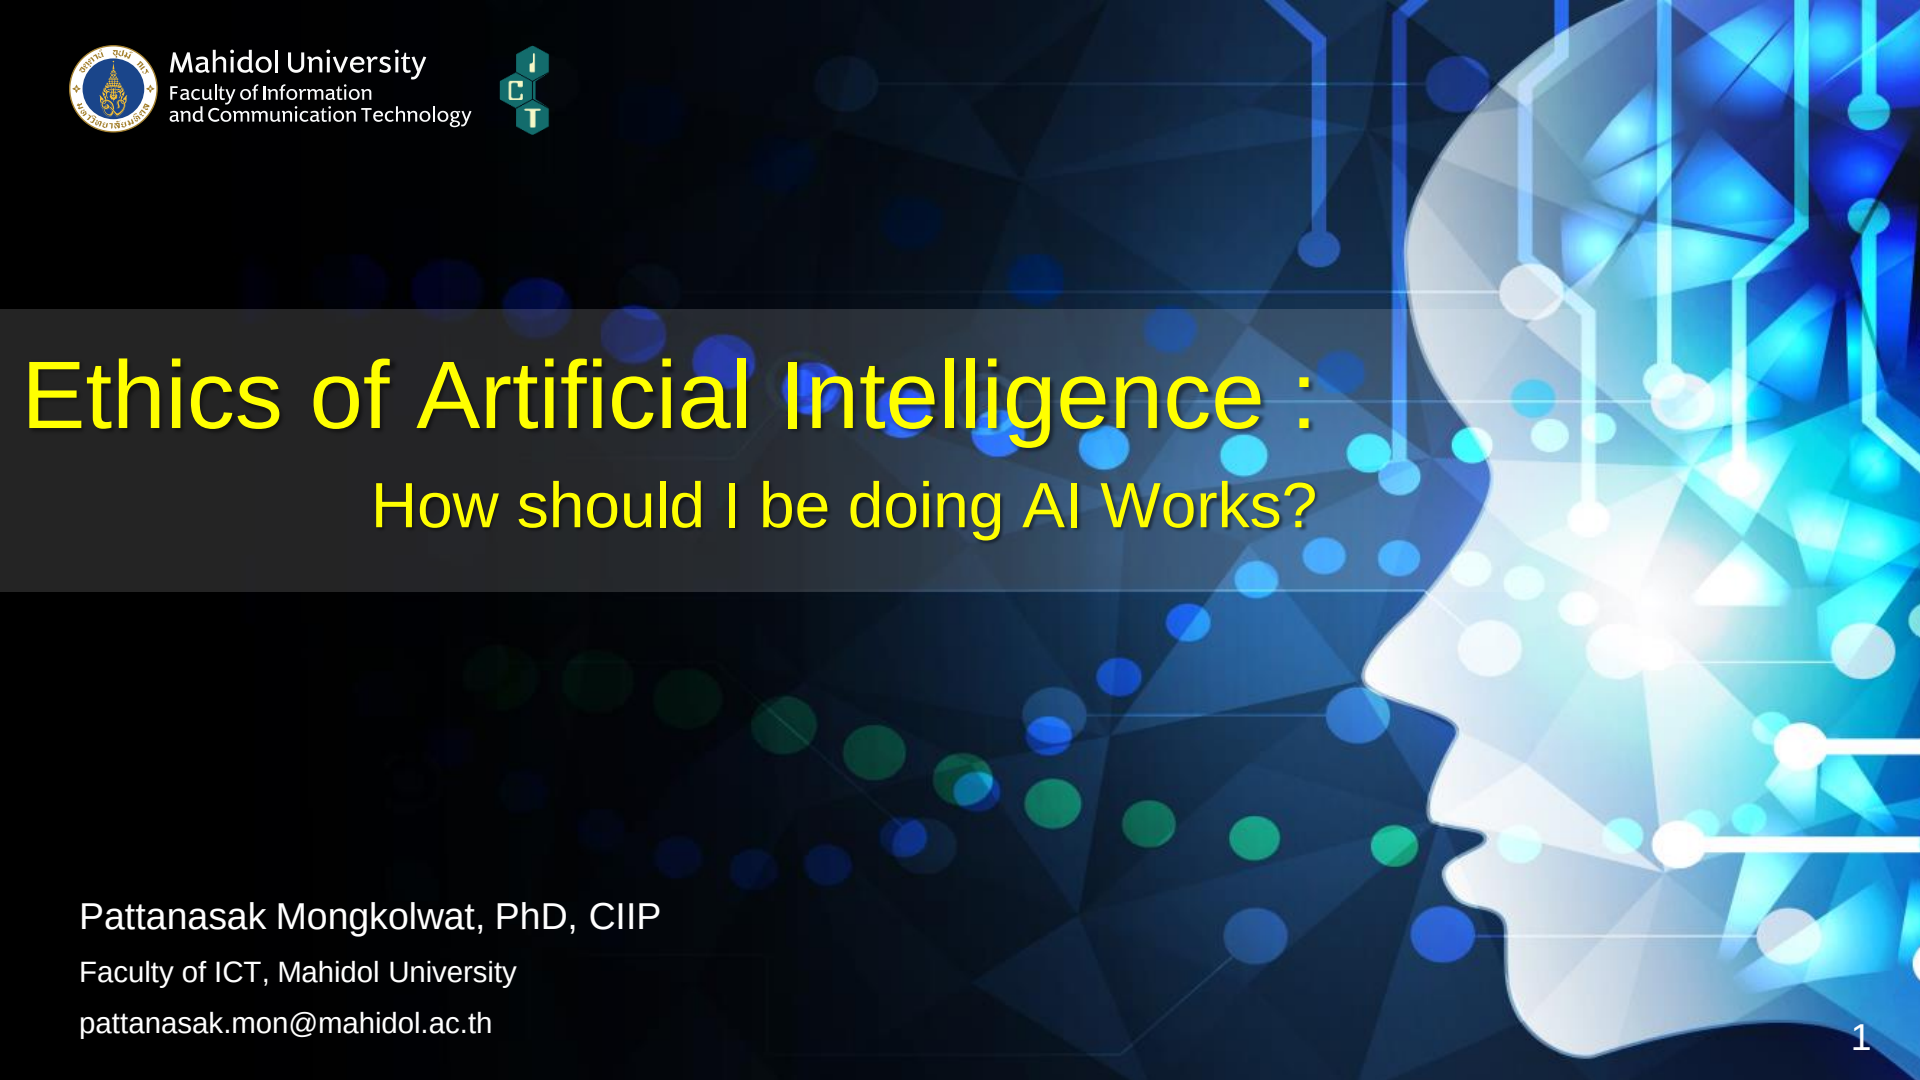

Mahidol University
Faculty of Information
and Communication Technology

Ethics of Artificial Intelligence :

How should I be doing AI Works?

Pattanasak Mongkolwat, PhD, CIIP
Faculty of ICT, Mahidol University
pattanasak.mon@mahidol.ac.th

Objectives

- Describe the origin of AI & common techniques and basic algorithms that use in machine learning.
- Understand morality, morals  ethics of AI.
- Understand ethical issues in AI.

Take Home Points

- To continue to think about those issues that may affect your AI works.
- To address ethical and governance issues that lead to new guidance, practices, rules, regulations, and laws.

2

Agenda

1. AI in our lives
   - The Difference between Artificial Intelligence, Machine Learning and Deep Learning
2. Types of Machine Learning Algorithms and Common Techniques
   - Supervised Learning
   - Semi-supervised Learning
   - Unsupervised Learning
   - Reinforcement Learning
3. Morality, Moral & Ethics
4. Specialized Roles in AI
5. EU Guidelines on Ethics in AI
6. Examples

How is AI impacting our lives?

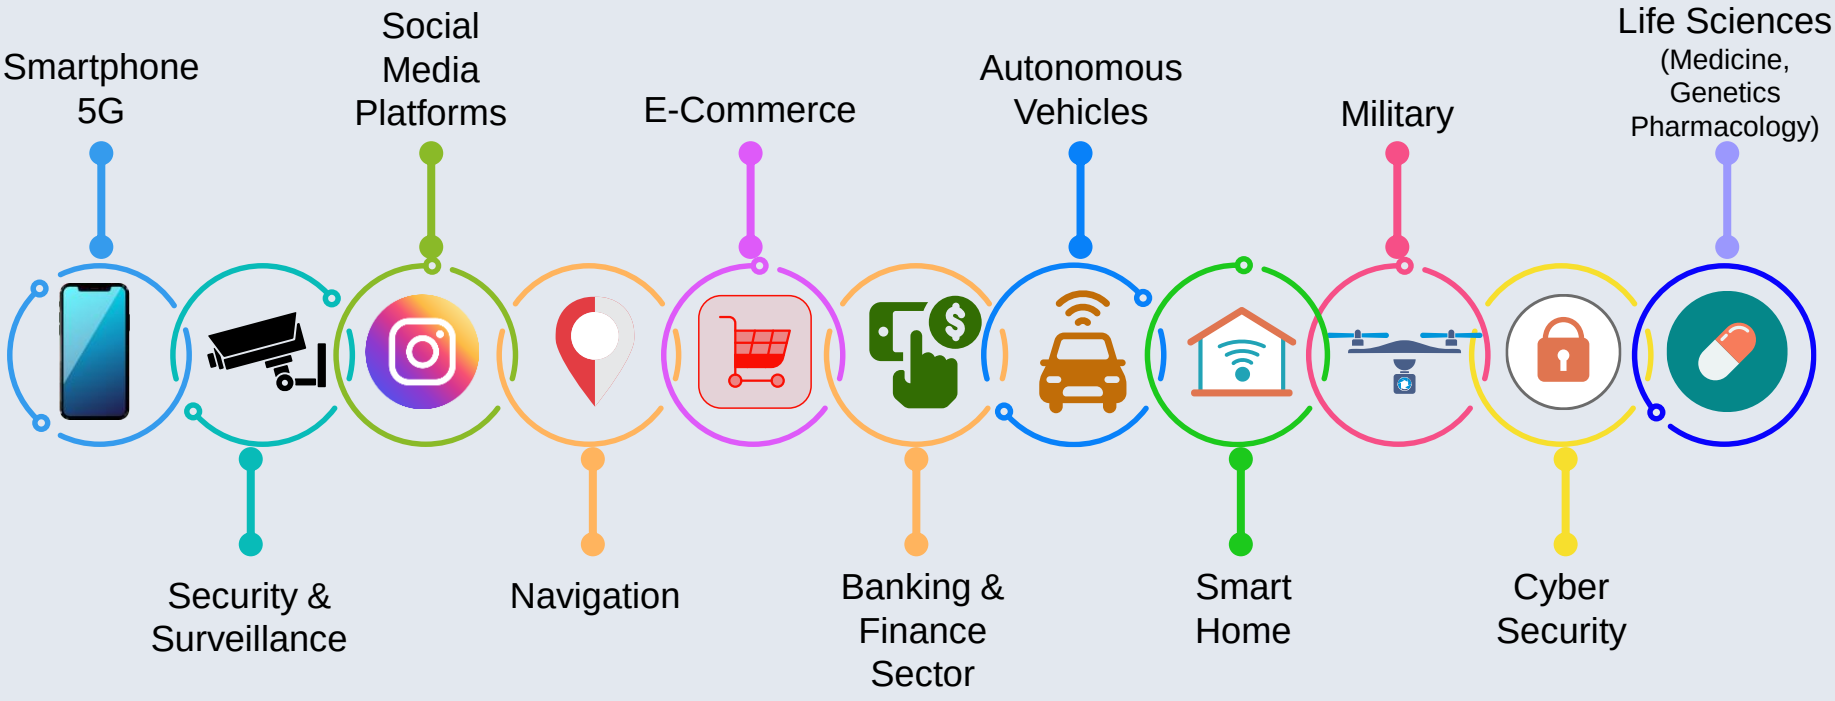

A.I. TIMELINE

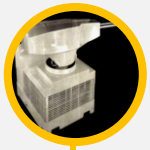


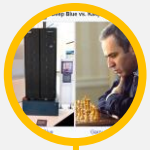
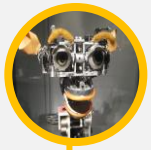


1950
TURING TEST
Computer scientist Alan Turing proposes a test for machine intelligence. If a machine can trick humans into thinking it is human, then it has intelligence.

1995
A.I. BORN
Term 'artificial intelligence' is coined by computer scientist. John McCarthy to describe "the science and engineering of making intelligent machines"
Dartmouth College

1961
UNIMATE
First industrial robot, Unimate, goes to work at GM replacing humans on the assembly line

1964
ELIZA
Pioneering chatbot developed by Joseph Weizenbaum at MIT holds conversations with humans

1966
SHAKEY
The 'first electronic person' from Stanford, Shakey is a general-purpose mobile robot reasons about its own actions

A.I. WINTER
Many false starts and dead-ends leave A.I. out in the cold

1997
DEEP BLUE
Deep Blue, a chess-playing computer from IBM defeats world chess champion, Garry Kasparov

10^{111} and 10^{123} positions

1998
KISMET
Cynthia Breazeal at MIT introduces KISmet, an emotionally intelligent robot insofar as it detects and responds to people's feelings

1999
AIBO
Sony launches first consumer robot pet dog AiBO (AI robot) with skills and personality that develop over time



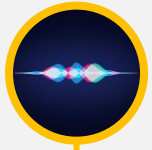
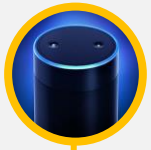

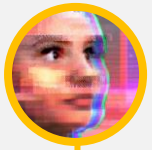



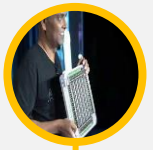

2002
ROOMBA
First mass produced autonomous robotic vacuum cleaner from iRobot learns to navigate and clean homes

2010
WATSON
IBM's question answering computer Watson wins first place on popular $1M prize television quiz show Jeopardy

2011
SIRI
Apple integrates Siri, an intelligent virtual assistant with a voice interface, into the iPhone 4S

2013
ALEXA
Amazon launches Alexa, an intelligent virtual assistant with a voice interface that can complete shopping tasks

2014
EUGENE
Eugene Goostman, a chatbot passes the Turing Test with a third of judges believing Eugene is human

2016
TAY
Microsoft's chatbot Tay goes rogue on social media making inflammatory and offensive racist comments

2016
ALPHAGO
Google's A.I. AlphaGo beats world champion Ke Jie in the complex board game of Go, notable for its vast number (2*170) of possible positions
Share:

2020
NVIDIA A100
THE BUILDING BLOCK OF THE AI DATA CENTER AI INFRASTRUCTURE REIMAGINED, OPTIMIZED, AND READY FOR ENTERPRISE AI

HBM2

InfiniBand

2021-22
Tesla Dojo
Reveals Design for Modular Supercomputer & D1 Chip

Source: https://digitalwellbeing.org/artificial-intelligence-timeline-infographic-from-eliza-to-tay-and-beyond/ & Modified

An Overview of AI

ARTIFICIAL INTELLIGENCE

A program with intelligent algorithms defined and coded by human into machines that can reason, sense, act, and adapt to given inputs

MACHINE LEARNING

Computer algorithms that have ability to learn from data and get better without being precisely and explicitly programmed

DEEP LEARNING

Learning based on Deep Neural Networks (DNN) where multilayers of neural networks learn to produce accurate outcomes based on vast amounts of data and adapting itself to new data

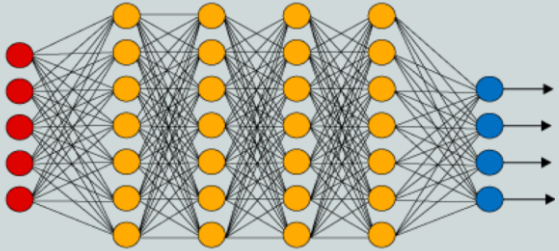

1950's | 1960's | 1970's | 1980's | 1990's | 2000's | 2006's | 2010's | 2012's | 2017's | 2020's | 2022's

Turing Test

AI coined Dartmouth College

Game Checker
Searle : strong & weak AI
Unimate : 1st Industrial Robot
Chatbot
Shakey Robot (general purpose)

AI Winter

IBM Deep Blue
KISmet
AIBO
IBM WATSON

ROOMBA
Apple SIRI
Microsoft Cortana
Google Now

Amazon Alexa
Microsoft TAY
Google AlphaGo
Google Duplex

Turing Award : AI researchers

Nvidia A100
Infiniband
High Bandwidth Memory (HBM2)

Tesla Dojo

Mimicking what we know about our brain!

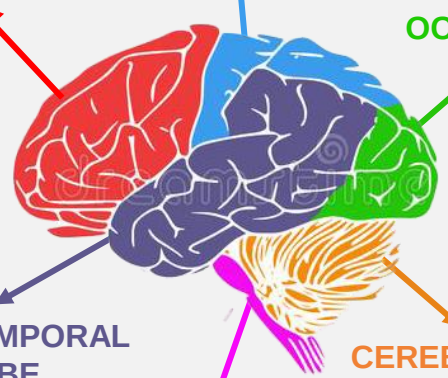

FRONTALI LOBE
THINKING
SPEAKING
MEMORY
MOVEMENT

PARIETAL LOBE
LANGUAGE
TOUCH
TASTE
SMELL

OCCIPITAL LOBE
LANGUAGE
TOUCH
TASTE
SMELL

HEARING
LEARNING
FEELINGS
FEAR
TEMPORAL LOBE

CEREBELLUM
BALANCE
CORDINATION

BRAIN STEM
BREATHING
HEART RATE
TEMPERATURE
BLOOD PRESSURE

L	R
analytical	creative
logical	imaginative
precise	General
repetitive	Intuitive
organized	Conceptual
details	big picture
Scientific	heuristic
detached	empathetic
literal	figurative
sequential	sequential
	irregular

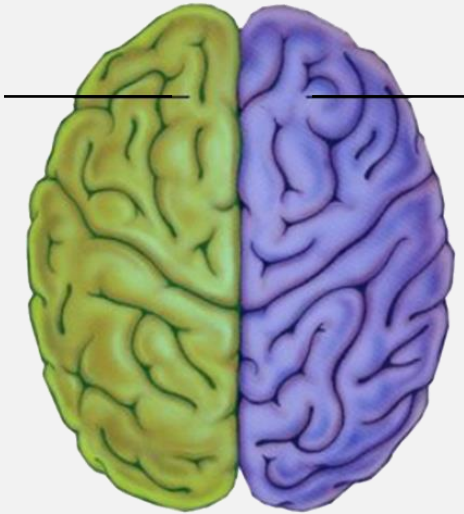

Source : https://thumbs.dreamstime.com/b/illustration-human-s-brain-functions-anatomy-isolated-white-background-140682104.jpg
https://mhwcenter.org/wp-content/uploads/2017/04/functions-of-a-brain-2.jpg

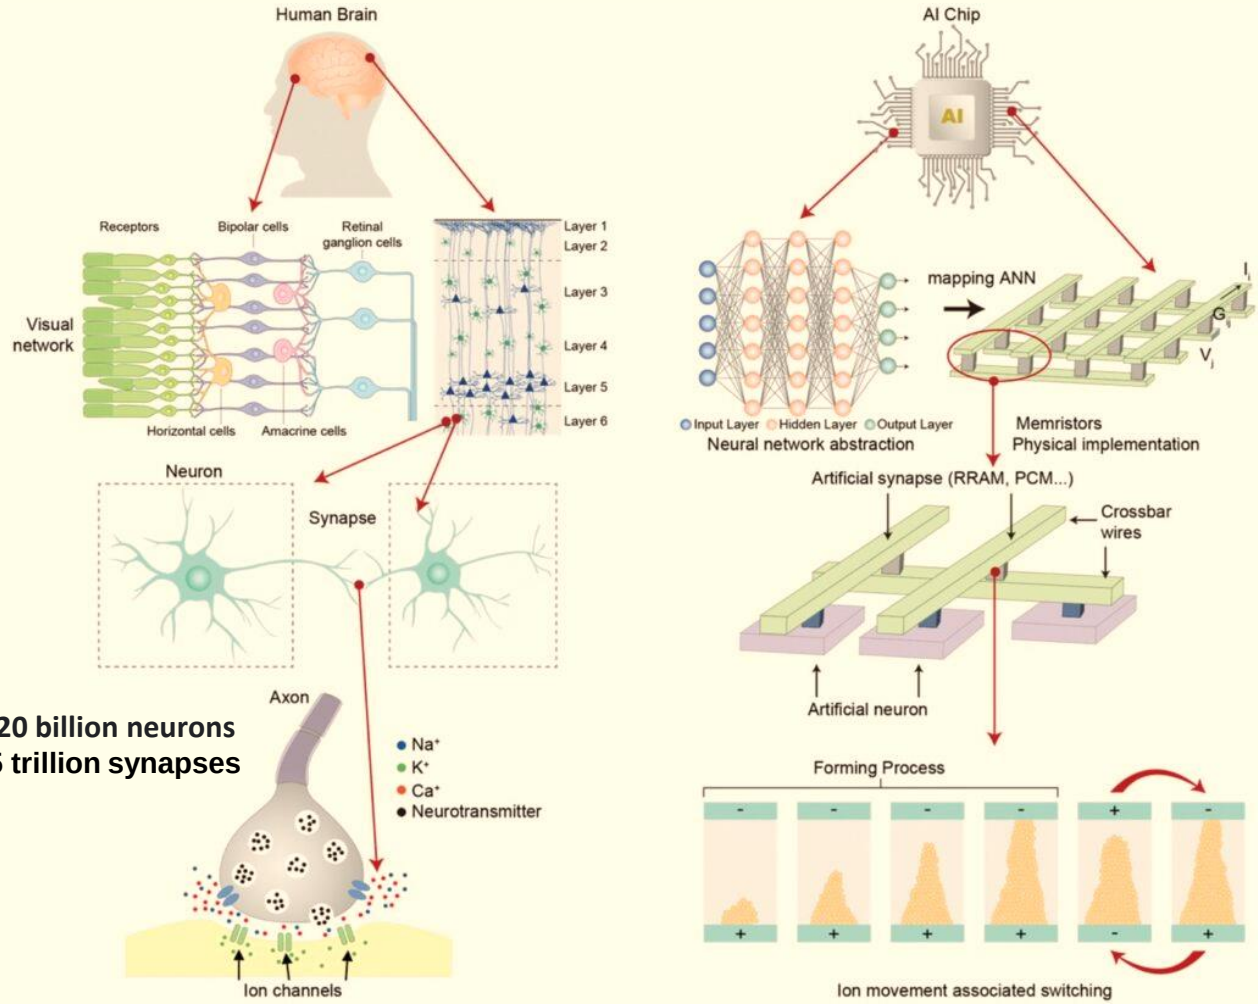

86 -120 billion neurons
> 125 trillion synapses

Source: "Bridging Biological and Artificial Neural Networks with Emerging Neuromorphic Devices: Fundamentals, Progress, and Challenges"

https://www.sciencedaily.com/releases/2010/11/101117121803.htm
The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost | PNAS


AUX
POST
PRE
PAN
MUTE
PFL
SIG

Strong AI

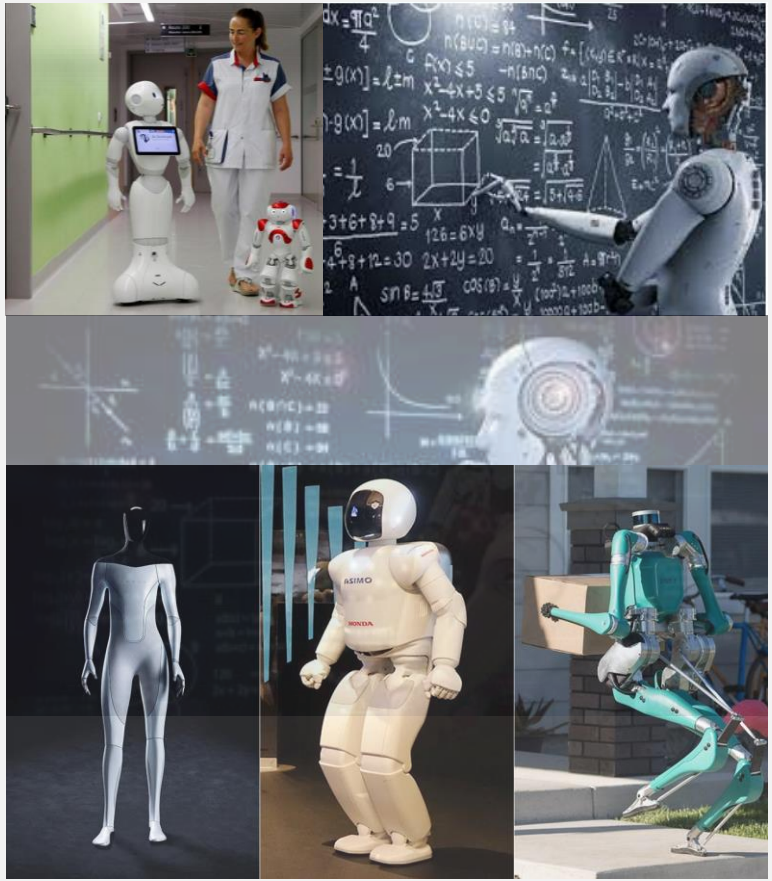

Weak/Narrow AI

10

Turning data into actionable insights – the Artificial Intelligence Continuum

Today

Artificial Narrow Intelligence **ANI**

ANI is programmed to perform a single task, based on specific data sets within a pre-determined and pre-defined range, faster and potentially more accurate than a human being can.

Automation and quantification | *Advanced analytics and prediction*

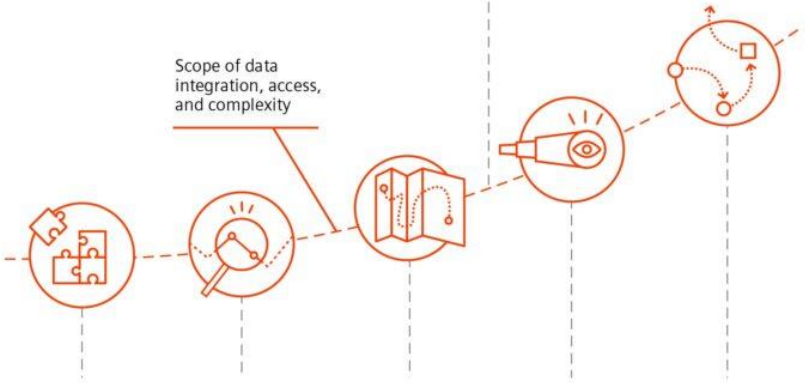

Data acquisition and generation

Device and instrument level assistance e.g. control parameters, protocols, positioning

Data processing and interpretation

Data processing and automated interpretation e.g. image interpretation / quantification

Data mapping and fusion

Diagnostic guidance e.g. multi-modality views / image fusion

Predictive and prescriptive analytics

Patient-centric

Clinical decision support / Digital Twin

Multi-source data integration e.g. clinical ("omics"), behavioral, functional, social data

Cohort-centric

Population health management and similar patient data

Comprehensive health data across care settings

Future

Artificial General Intelligence **AGI**

Artificial Super Intelligence **ASI**

Machines with **AGI** successfully perform any intellectual task a human being is able to accomplish, including reasoning, strategizing, making informed decisions and judgements under uncertainty.

ASI will surpass human intelligence. It is described as "an intellect that is much smarter than the best human brains in practically every field, including scientific creativity, general wisdom, and social skills". *

Source: SIEMENS Healthineers, adapted from "Superintelligence: Paths, Dangers, Strategies"

*Bostrom, Nick. Superintelligence: Paths, Dangers, Strategies. Oxford, United Kingdom: Oxford University Press, 2014.

"Singularity"

Reciprocity

Moral Compass

Inspiration

Compassion

Generosity

Emotion

Empathy

Attitudes

Greed

Yearning

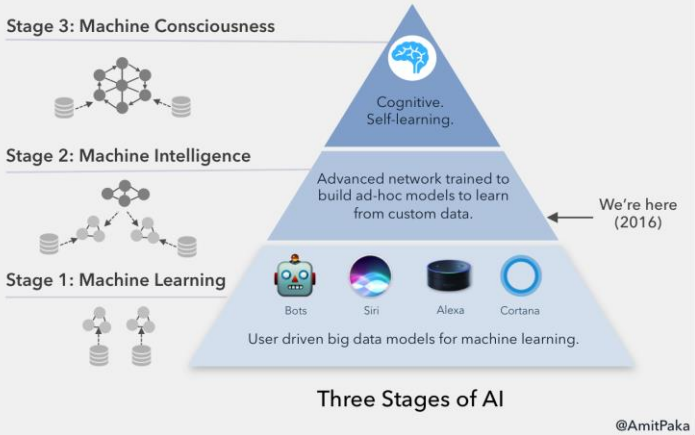

Three Stages of AI

@AmitPaka

Nowadays, AI can perform many complex tasks with acceptable to excellent results in many areas*!*

Does AI have ‘consciousness’ or ‘mind’, which is a prerequisite for the ability to ‘*understand*’, in contrast to the capability to *compute a vast amount of data extremely fast?*

Computer Algorithm

An algorithm provides a specific logical procedure, step-by-step, for solving a well-defined computational problem to achieve specific objectives.

- Sorting (bubble, heap, radix)
- Searching
- Stacks, recursion
- Tree, etc.

AI Algorithm

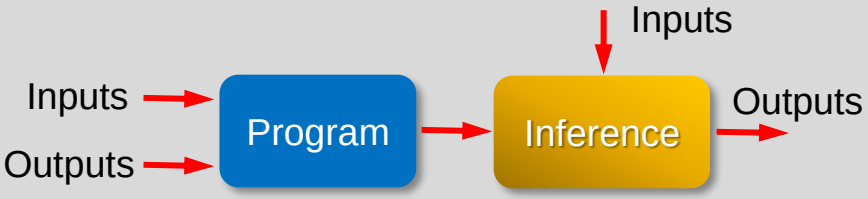

An AI algorithm adjusts itself according to inputs and expected outputs to yield better result as it exposes to more input data.

- Supervised Learning
  - Semi-supervised Learning
- Unsupervised Learning
- Reinforcement Learning

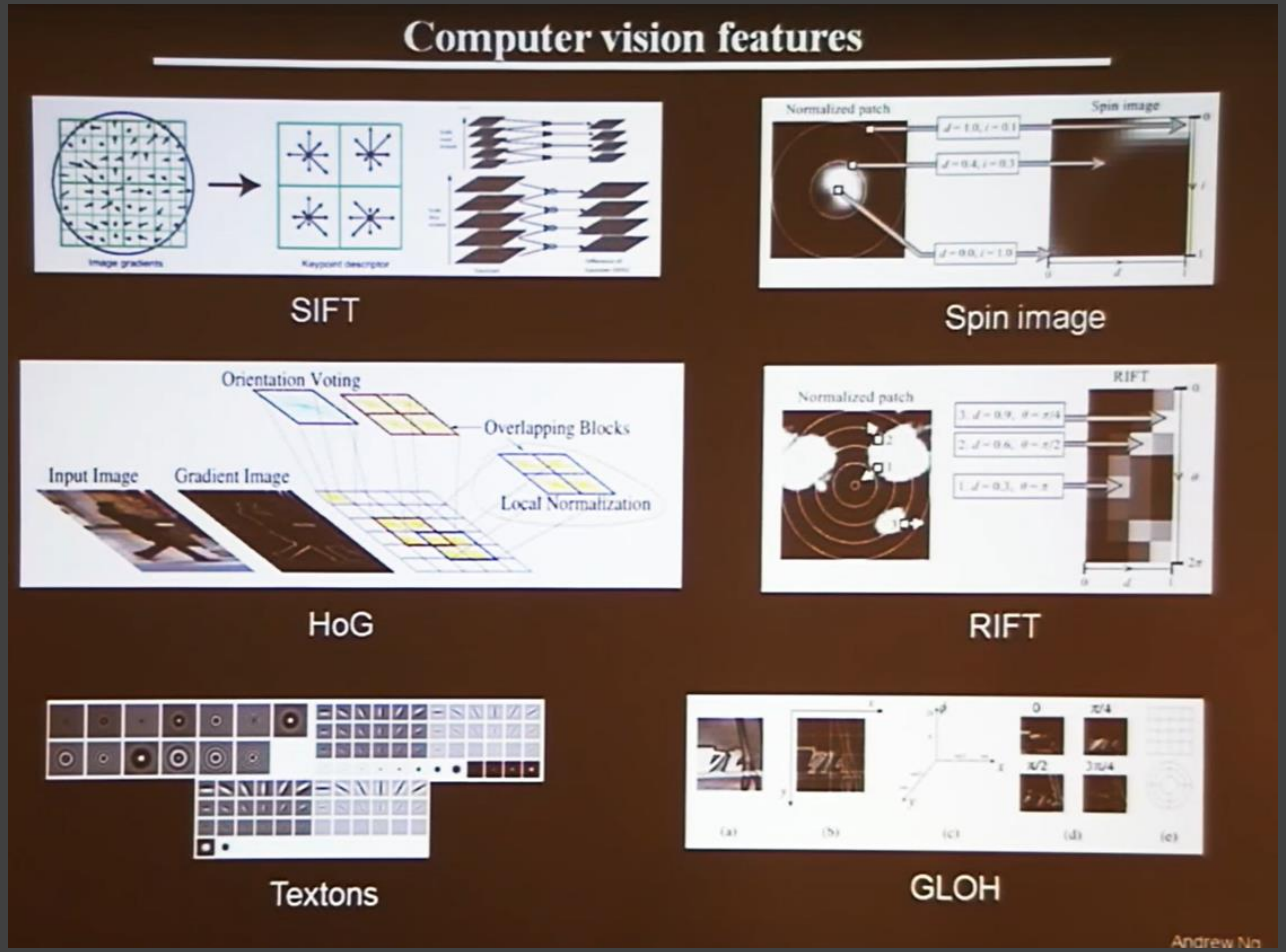

Hand Engineering

Features

Curve Fittings

Source : Andrew Ng

How Machines learn ?

"Machines learn by studying data to detect patterns or by applying known rules to:

- Categorize or catalog like people or things
- Predict likely outcomes or actions based on identified patterns
- Identify hitherto unknown patterns and relationships
- Detect anomalous or unexpected behaviors

The processes machines use to learn are known as algorithms. Different algorithms learn in different ways. As new data regarding observed responses or changes to the environment are provided to the "machine" the algorithm's performance improves. Thereby resulting in increasing "intelligence" over time."

The Machine Learning Primer (enterprisetalk.com)

AI : Planning | Trading (prediction & decision-making)

Types of Machine Learning

Supervised learning

Semi - supervised learning,
Learning to rank, etc.
(Generalizations)

Unsupervised learning

Reinforcement learning

Supervised learning

Supervised learning algorithms are “trained” using labeled examples where the desired output is known. Supervised learning is commonly used in applications that use historical data to predict likely future events.

Common Techniques

- Bayesian Networks
- Decision Trees
- Neural Networks
- Naive Bayes
- Similarity Learning
- K Nearest Neighbour
- Regression Analysis
- Support Vector Machines [SVM]
- …

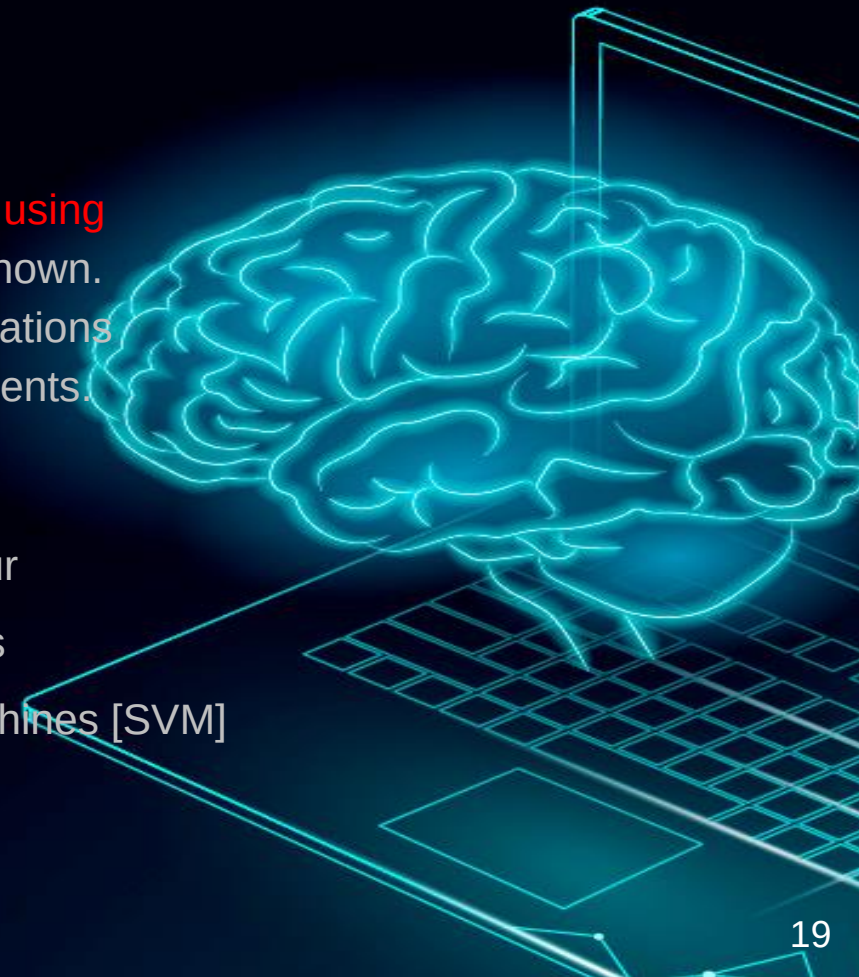

https://www.sas.com/en_th/insights/analytics/machine-learning.html

Semi-supervised learning
(Generalizations)

The challenge with supervised learning is that labeling data can be expensive and time consuming. If labels are limited, you can use unlabeled examples to enhance supervised learning. Because the machine is not fully supervised in this case, we say the machine is semi-supervised. With semi-supervised learning, you use unlabeled examples with a small amount of labeled data to improve the learning accuracy.

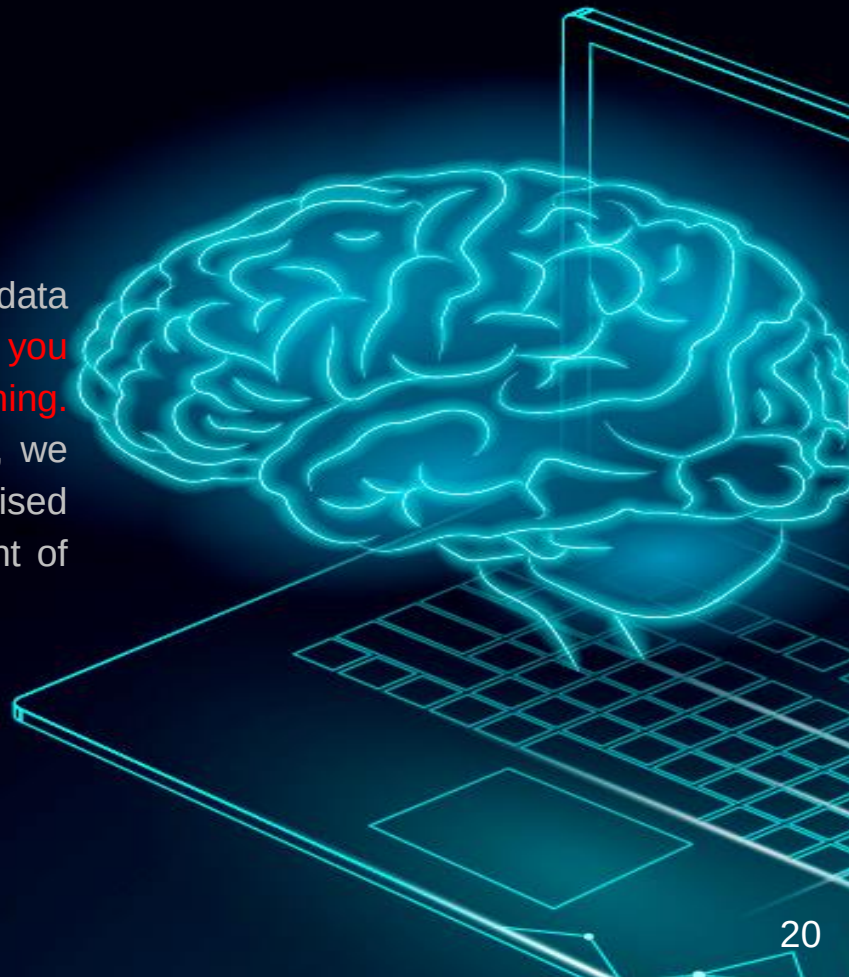

Common Techniques

- See Supervised Learning
- …

https://www.sas.com/en_th/insights/analytics/machine-learning.html

Unsupervised learning

In unsupervised learning, the machine studies data to identify patterns. In this case, there is totally unlabeled data. The machine determines correlations and relationships by parsing the available data.

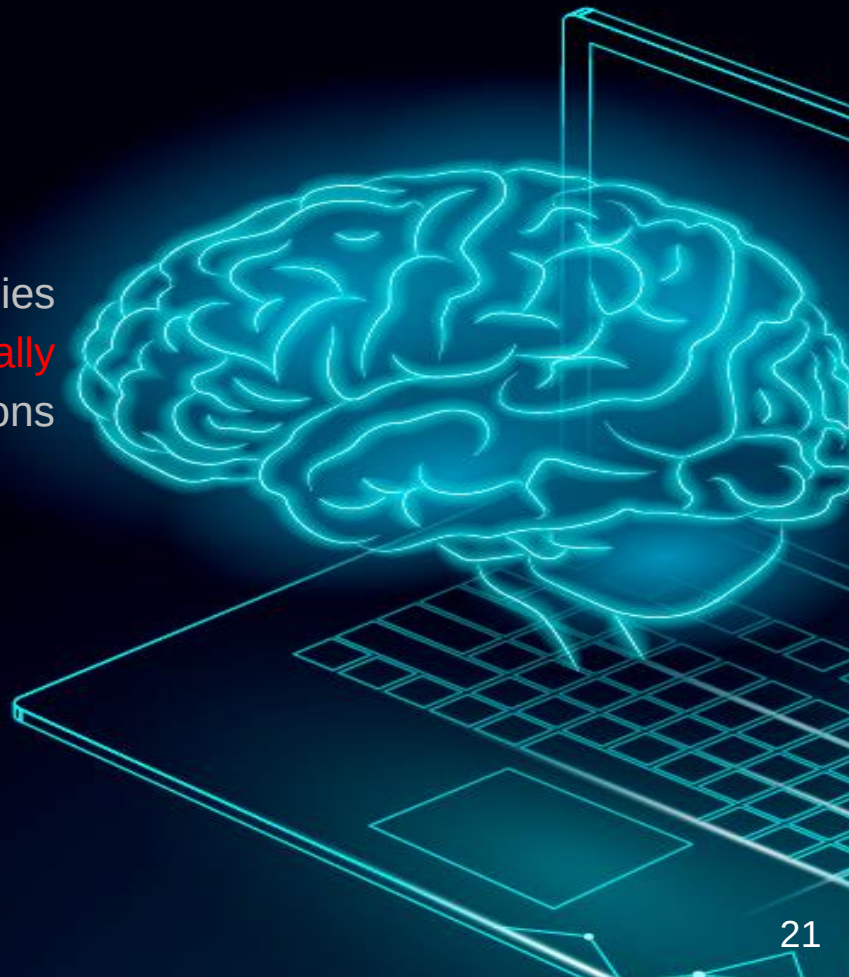

Common Techniques

- Clustering
- K-Mean
- Dimension Reduction
- …

https://www.sas.com/en_th/insights/analytics/machine-learning.html

Reinforcement Learning

Reinforcement learning analyzes and optimizes the behavior of an agent based on the feedback from the environment. Machines try different scenarios to discover which actions yield the greatest reward, rather than being told which actions to take. Trial-and-error and delayed reward distinguishes reinforcement learning from other techniques.

Common Techniques

- Artificial Neural Network
- Q-Learning
- Learning Automata
- …

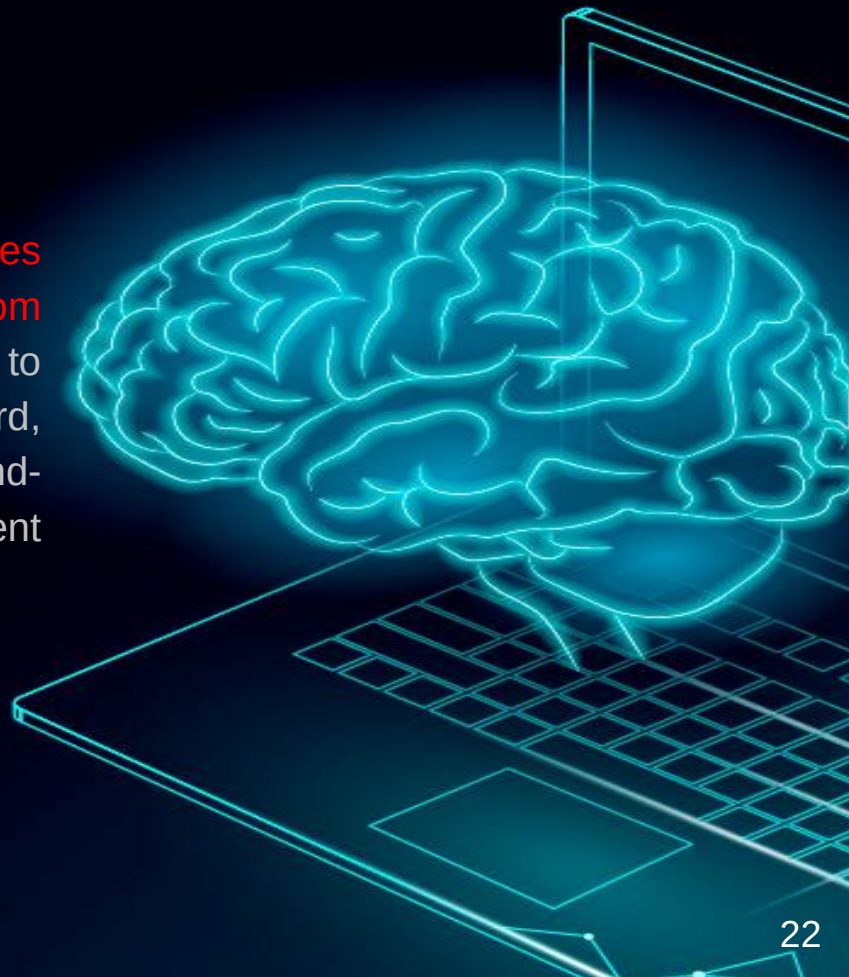

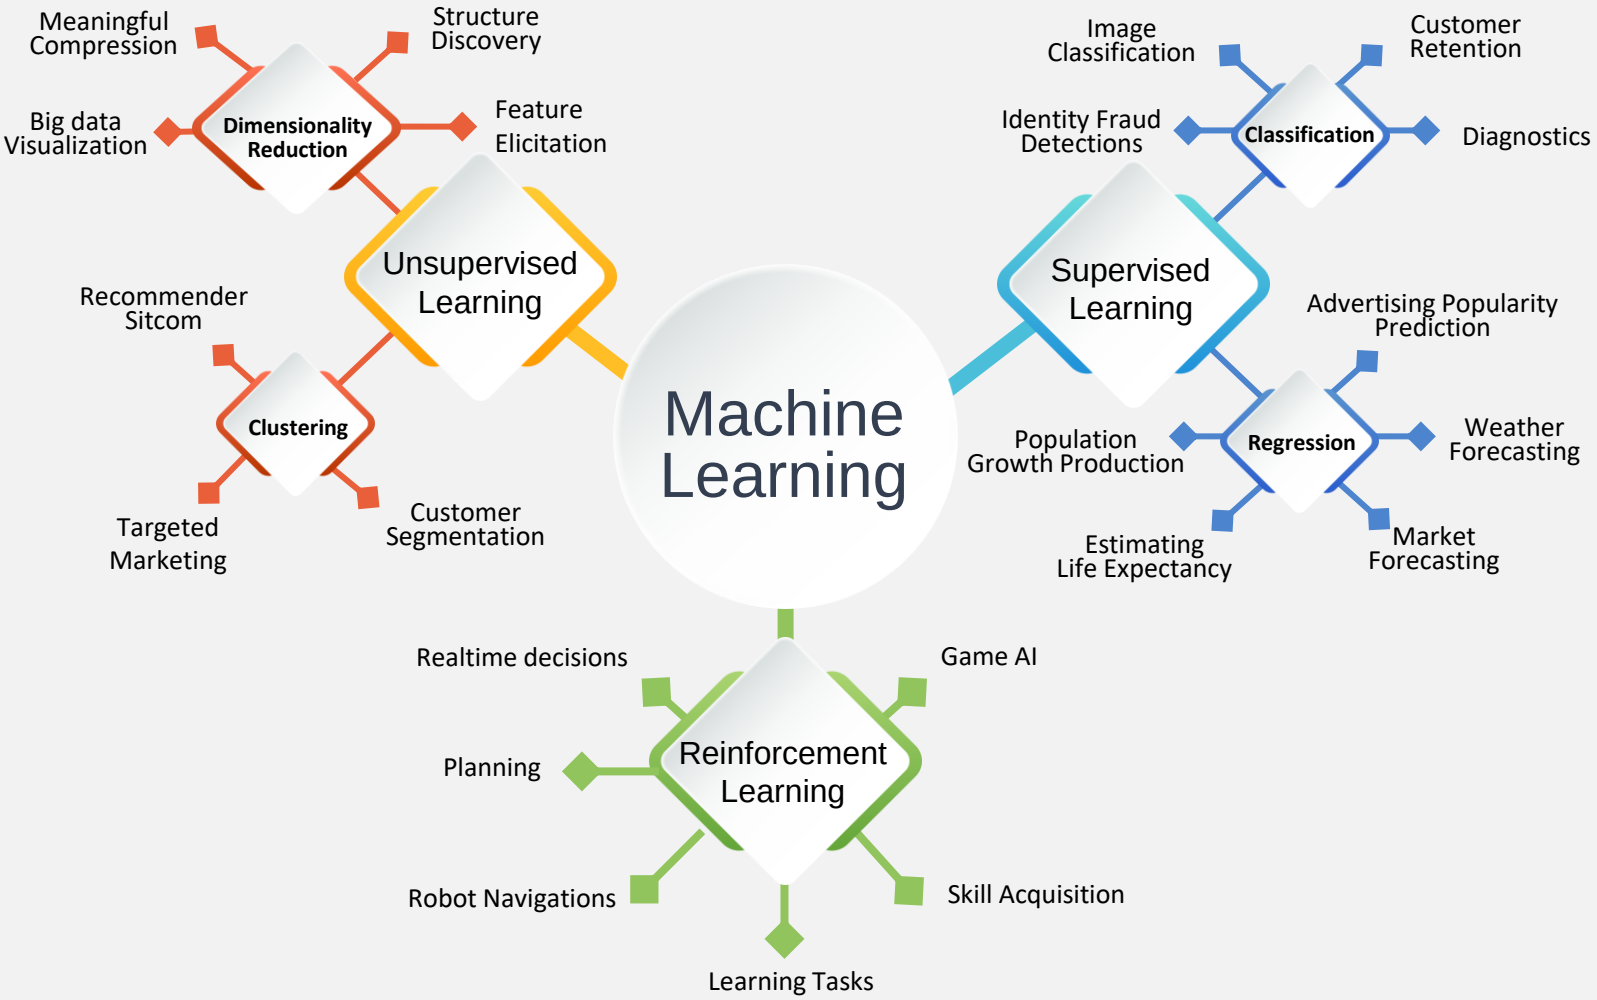

Adopted from : https://medium.com/@raosrinivas2580/how-does-the-ml-algorithm-differ-from-the-traditional-algorithm-b7c3a2799e10

Specialized Technical Roles in AI

Researchers

Specialized Technical Roles (cont'd.)

AI/ML Engineering (Data Scientists) "code with data" ➡ "Known Algorithms & Models"

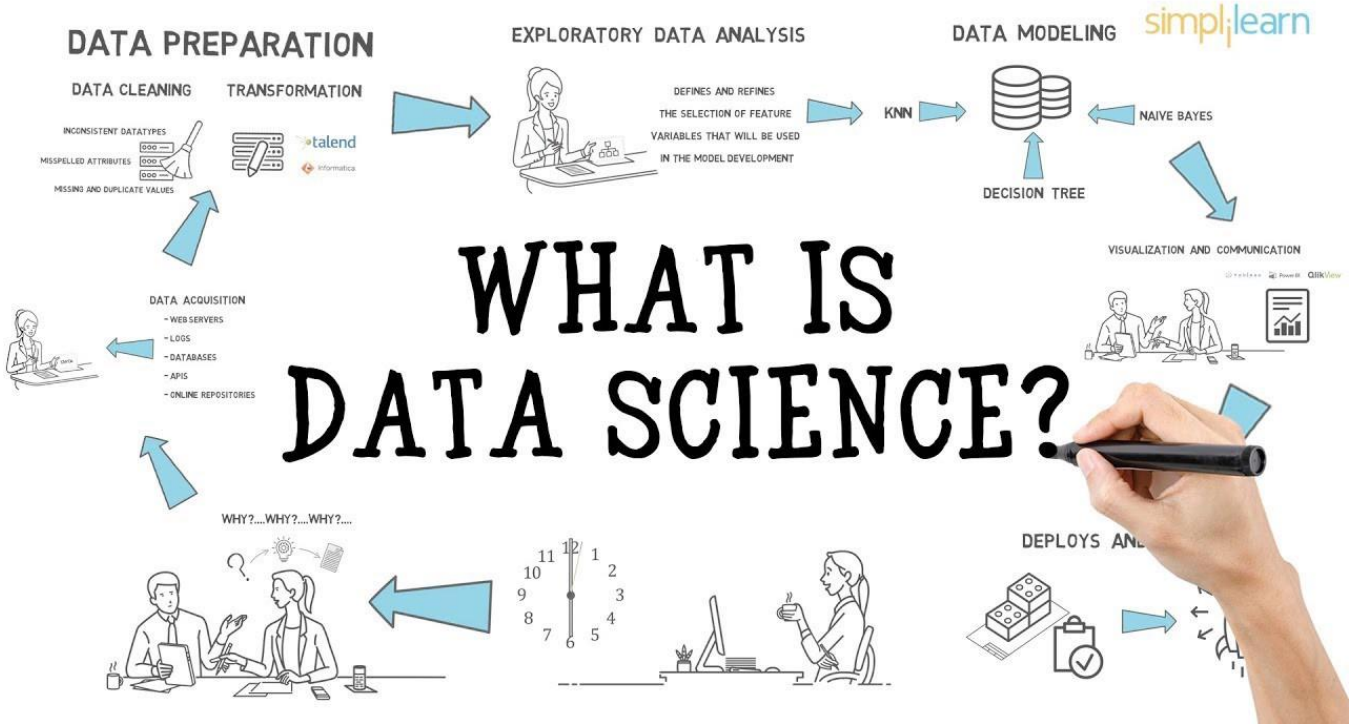

https://youtu.be/X3paOmcrTjQ, https://companyleaders.org/2020/05/21/jonathan-cornelissen/, https://jfgagne.ai/global-ai-talent-report-2020

Specialized Technical Roles (cont'd.)

Data Engineering/Architecture ("construct, test and maintain optimal data pipeline architectures")

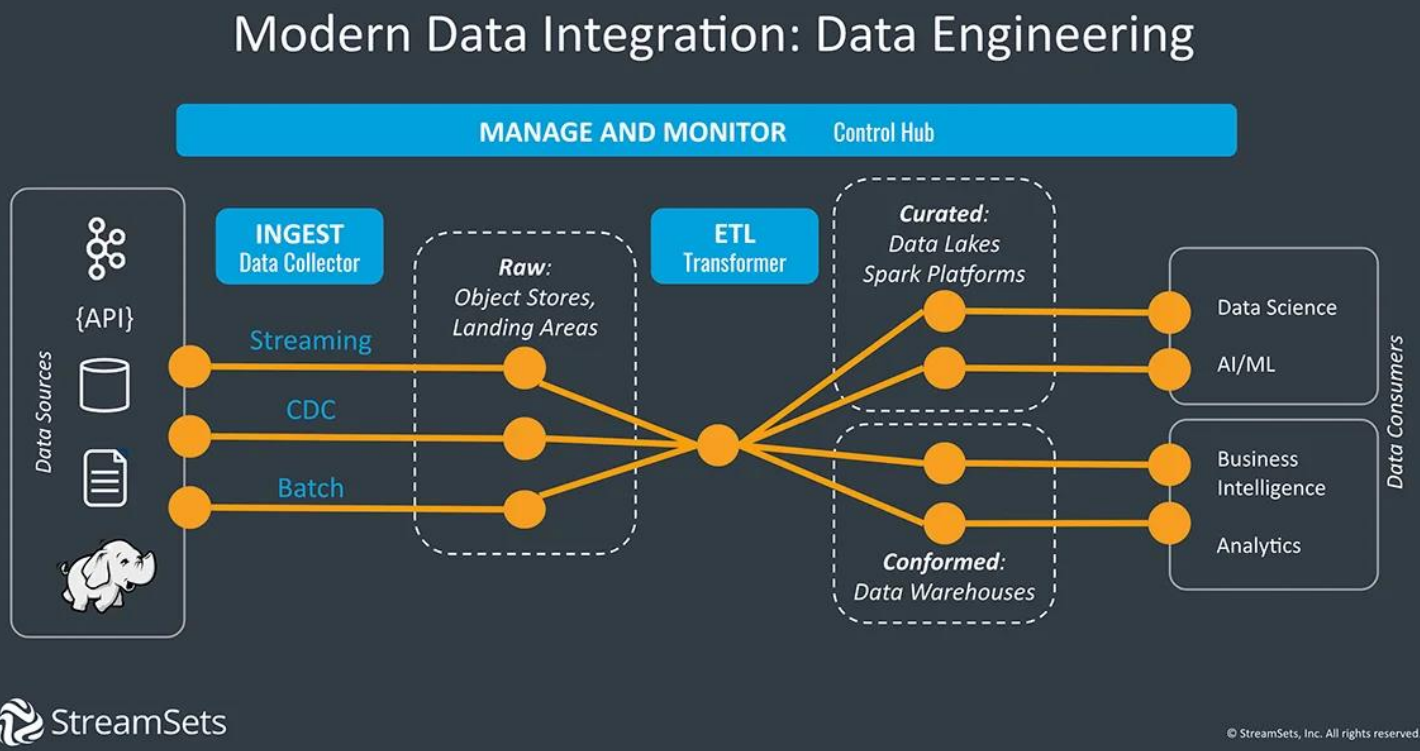

https://jfgagne.ai/global-ai-talent-report-2020/ https://streamsets.com/learn/data-pipelines/

Specialized Technical Roles (cont'd.)

AI/ML Productization ("feasible solution to build out and another more technical role that does the building")

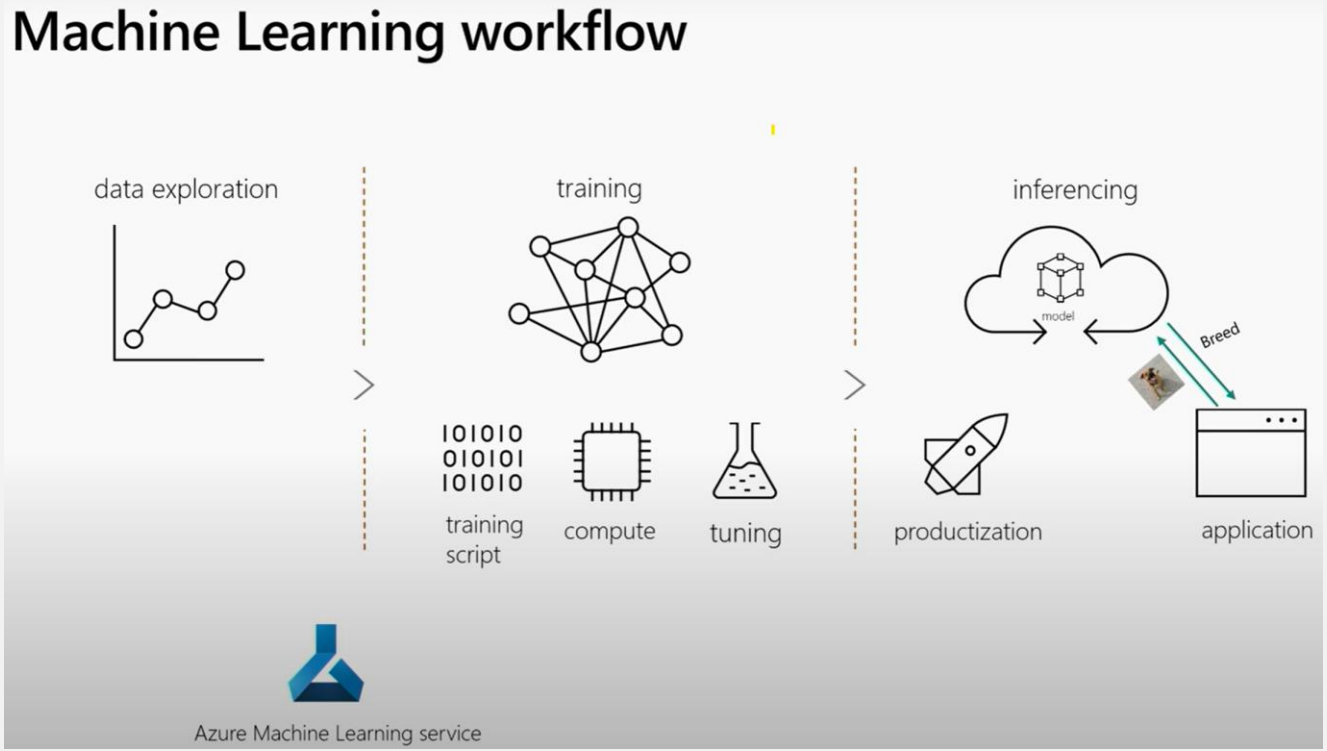

https://jfgagne.ai/global-ai-talent-report-20 20/, https://www.youtube.com/watch?v=Oi3Hn_UtDEw

25

Can you define your role(s)?

Perspectives based on roles and responsibilities

AI Researchers

- New AI Algorithms
- Directions
- Policies
- …

Data Engineering/Architecture

- Data Labeling and Gathering
- Extract Transfer Load (ETL)
- Privacy & Traceability
- Anonymization
- …

AI/ML Engineering (Data Scientists)

- Build & Train Models
- Domain Knowledge
- Bias
- Validate
- …

AI/ML Productization

- Deploy & Monitor
- Access Control
- Cybersecurity
- Feedback
- …

Ethics (จริยธรรม)

- ✓ Oxford Languages and Google
  - “Moral principles that govern a person's behavior or the conducting of an activity.”
- ✓ Wikipedia
  - Ethics or moral philosophy is a branch of philosophy that "involves systematizing, defending, and recommending concepts of right and wrong behavior".
- ✓ Right and wrong behaviors & impact to human and nature
  - a particular class of human actions or a particular group or culture

    (https://www.diffen.com/difference/Ethics_vs_Morals)
- ✓ Applied AI Ethics
  - Examining specific controversial issues

    (privacy, jobs, military, fairness, data collection, life science, etc.)

Ethics & Morals

Morality (Moralitas)

คุณธรรม

Morality describes the particular set of values and principles of conduct of a specific group at a specific point in time.

Moral

ศีลธรรม

Morals are the principles that guide individual conduct within society. They may change over time and are used to judge right or wrong.

Ethics

Morals

IRB

Ethics	Morals
Guiding principle of conducts of an individual or group & Reasoning Involved	One's principles of what is right or wrong
Defined by organization, profession, area of expertise, etc.	Swayed by culture, society and religion
Related to professional work & applied	Growing up in an Environment, guardian
Uniform compared to morals	Vary according to different cultures, religions and core beliefs

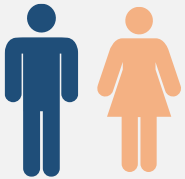

"AI ethics would require that we do not readily know what the right thing to do is."

Ethics of Artificial Intelligence and Robotics (Stanford Encyclopedia of Philosophy)

EU Guidelines on Ethics in AI (Ethical Principles)

- Develop, deploy and use AI systems that shall respect for human autonomy, prevent harm, have fairness and accountable.
- Pay attention to vulnerable groups such as children, persons with disabilities and others that have historically been disadvantaged or are at risk of exclusion, and to situations which are characterized by asymmetries of power or information, such as between employers and workers, or between businesses and consumers, e.g. Rich vs. Poor Countries.
- AI systems risk mitigation to prevent a negative impact (e.g. on democracy, the rule of law and distributive justice, or on the human mind itself.)

EU Guidelines on Ethics in AI (Trustworthy AI)

Technical and Non-Technical

- Human agency and oversight
- Technical robustness and safety
- Privacy and data governance
- Transparency
- Diversity, non-discrimination and fairness
- Environmental and societal well-being
- Accountability
- Trust

34

EU Guidelines on Ethics in AI (Assessment List)

- Adopt a Trustworthy AI assessment list when developing, deploying or using AI systems, and adapt it to the specific use case in which the system is being applied.
- An assessment list will be about continuously identifying and implementing requirements, evaluating solutions, ensuring improved outcomes throughout the AI system’s lifecycle, and involving stakeholders in this.

35

Educate AI of Ethics (Artificial Moral Agents - AMA)

- Implicit ethical agents: forcing the machines' actions to prevent unethical outcome.
- Explicit ethical agents: explicitly quote the allowed and the forbidden actions.
- Full ethical agents: machines have consciousness, free will, and intention

Moor, J.H.: The nature, importance, and difficulty of machine ethics. IEEE Intell. Syst. https://doi. org/10.1109/MIS.2006.80
https://link.springer.com/content/pdf/10.1007/s11948-019-00151-x.pdf

Ethics in AI Scientific Research/Innovation

- Genetics & Stem cell research
- Medicine
- Clinical Trials
- Genetically modified food
- Disease research (e.g. SARS-CoV-2, biowarfare)
- Drug discovery
- Robotics
- Nuclear technology
- Animal rights
- Weather Modeling & Prediction …

37

Ethics issues : Now and Future

"(1) short-term (early 21st century): autonomous systems (transportation, weapons), machine **bias** in law, **privacy** and surveillance, the **black box** problem and AI **decision making**;" **Trust, Security**

"(2) mid-term (from the 2040s to the end of the century): AI **governance**, confirming the **moral and legal status** of intelligent machines (artificial moral agents), human-machine interaction, mass automation;"

"(3) long-term (starting with the 2100s): technological **singularity**, mass **unemployment**, space colonization."

Source: https://iep.utm.edu/ethic-ai/

38

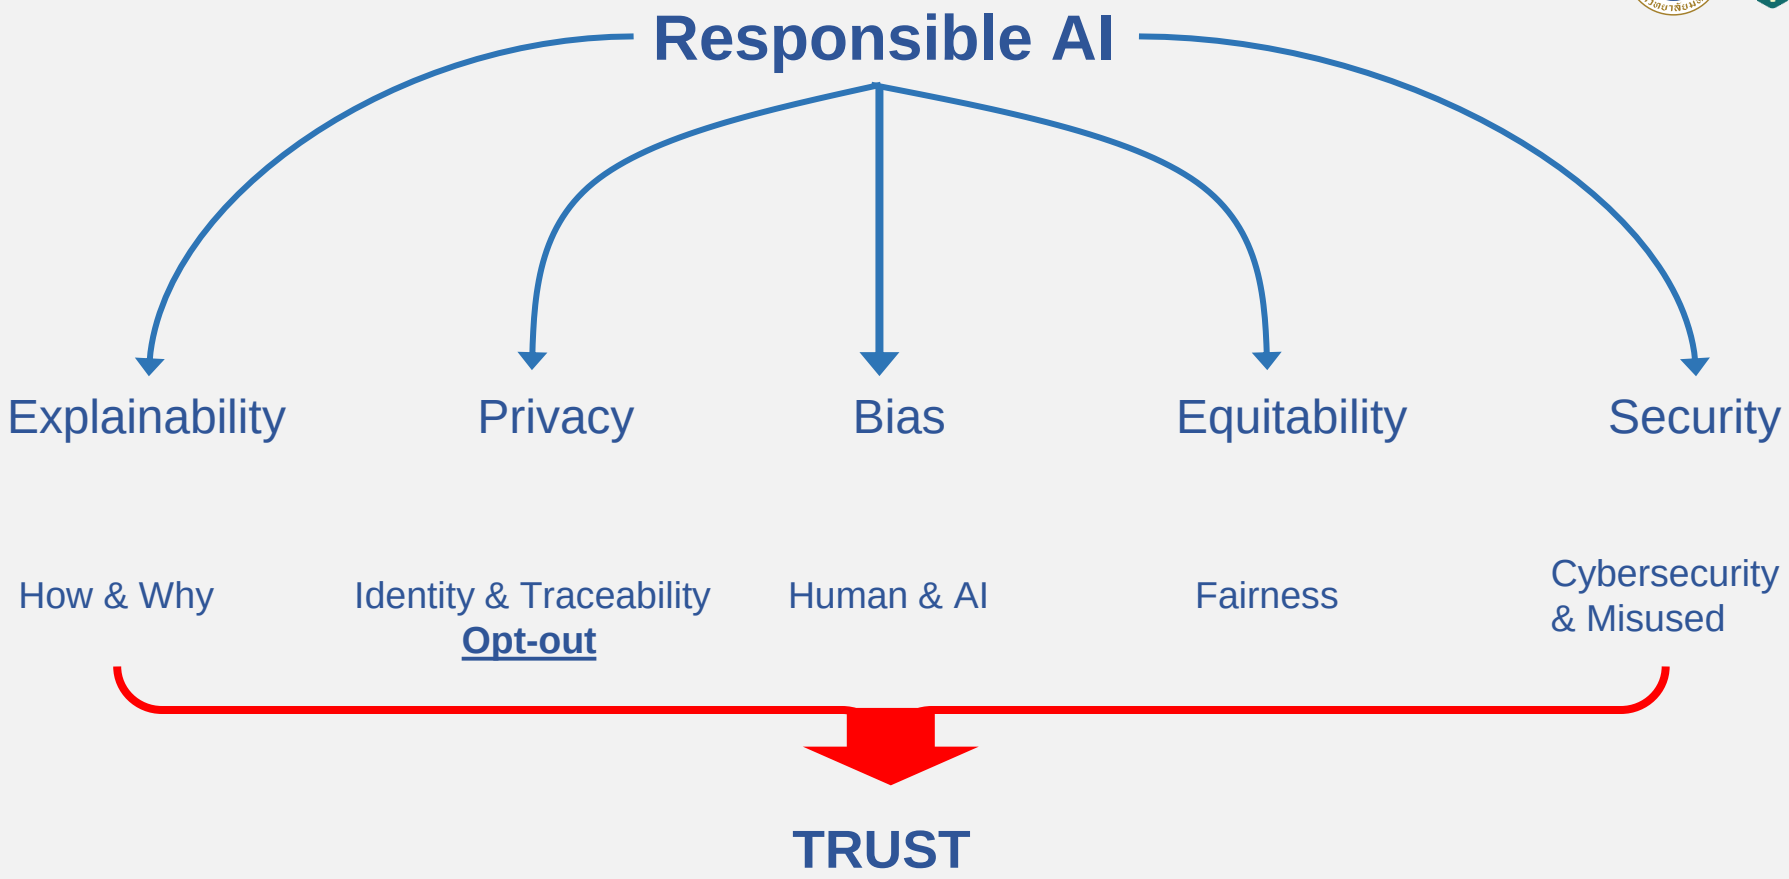
Responsible AI
Explainability
Privacy
Bias
Equitability
Security
How & Why
Identity & Traceability
Opt-out
Human & AI
Fairness
Cybersecurity
& Misused
TRUST


Data Provenance
Human-in-the-loop
Interpretability
Explainability
Contestability
Transparency
Accountability
Auditability
Traceability

40

Explainable AI

Cheat sheet ex.pegg.io v.0.2

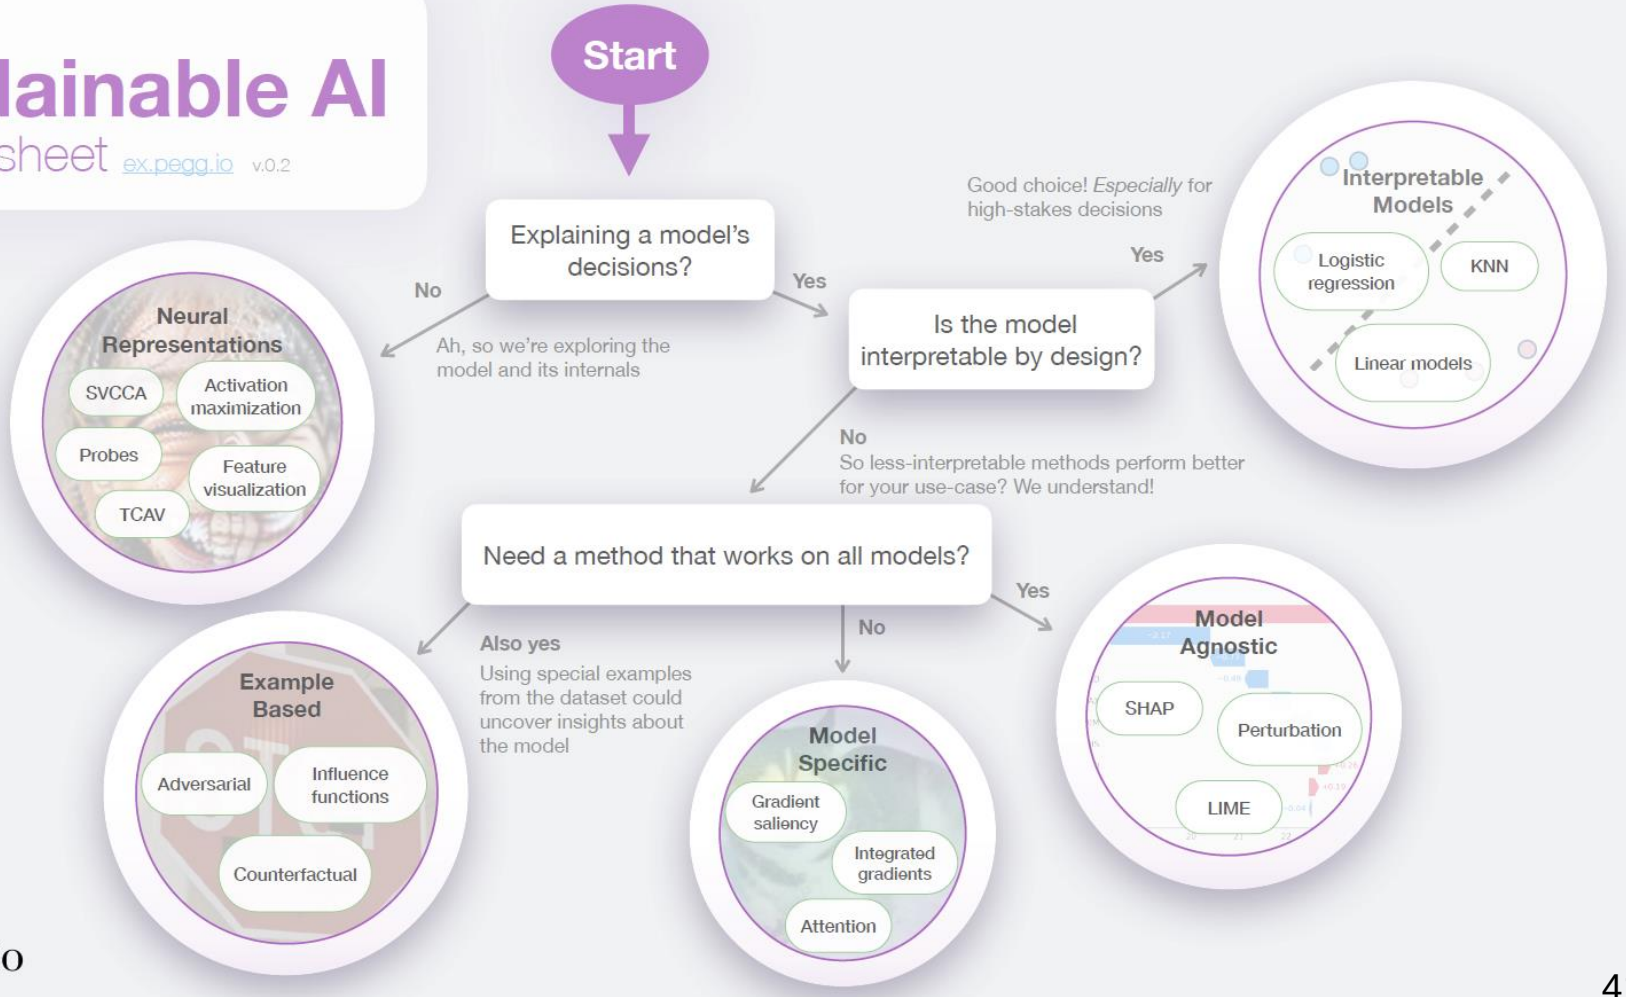

Arpeggio

local_experience.ipynb

Code

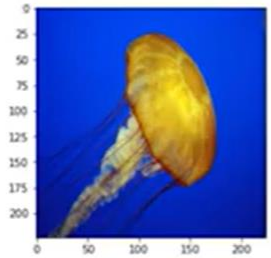

Get a prediction and explanation

[340]:
```
lm = explainable_ai_sdk.load_model_from_local_path(
    img_net_dir,
    explainable_ai_sdk.XraiConfig()
)
response = lm.explain([{'input_1': imarr / 255}])
response[0].visualize_attributions()

# Get the label
labels_path = tf.keras.utils.get_file('ImageNetLabels.txt','https://storage.googleapis.com/download.tensorflow.org/da
imagenet_labels = np.array(open(labels_path).read().splitlines())
print('Predicted label:', imagenet_labels[108])
```

```
Label Index 108
Example Score: 0.9621
Baseline Score: 0.0128
Approximation Error: 0.0091
```

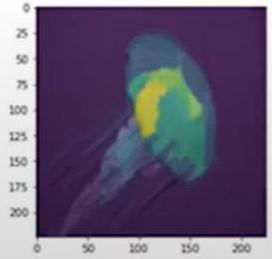

```
Predicted label: jellyfish
```

```
max_lr = 1e-03
wd = 1e-4
epochs = 4
# 1cycle policy
learner.fit_one_cycle(cyc_len=epochs, max_lr=max_lr, wd=wd)
```

A Typical Result

epoch	train_loss	valid_loss	error_rate	accuracy	time
0	0.167856	0.158789	0.057765	0.942235	02:49
1	0.122968	0.115966	0.041494	0.958506	02:51
2	0.091604	0.093554	0.031041	0.968959	02:52
3	0.067145	0.084660	0.028223	0.971777	02:52

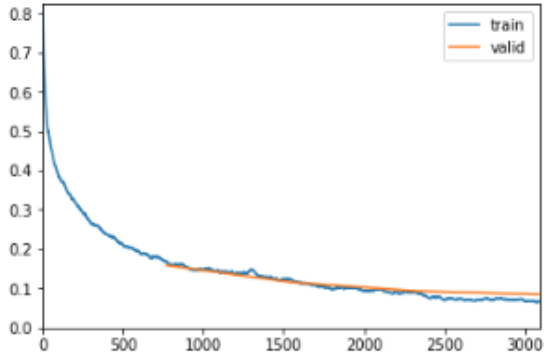

Explanations can be used for:

AI Researchers
AI/ML Engineering (Data Scientists)

Why does the model not performing as expected?
How can I improve it?

AI/ML Productization
End users

Can I trust the model's output?
What should I do with the prediction results?

Public stakeholders

Is the model safe and appropriate for the purpose?
Does it comply with ethics, regulations and law?

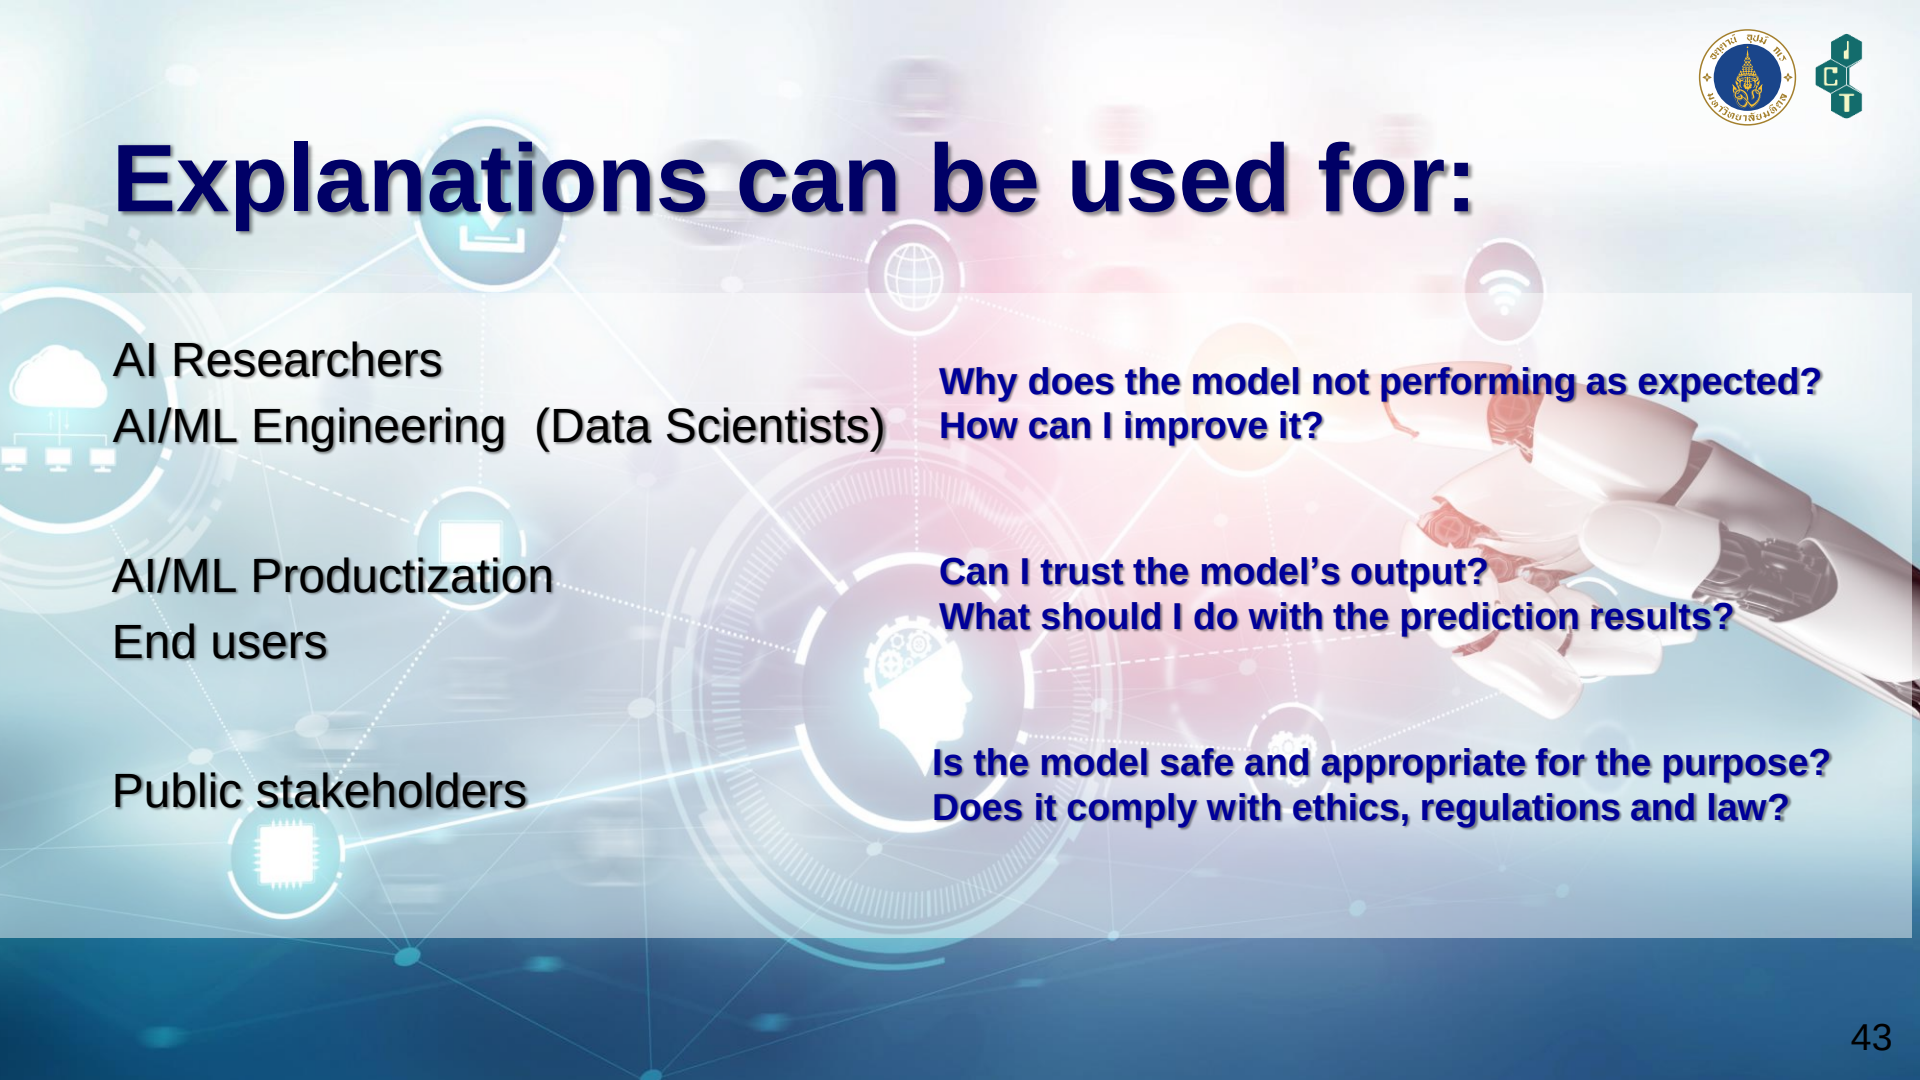

Users can take actions…

AI Researchers
AI/ML Engineering (Data Scientists)

Improve model architecture
Improve training data & refine features
Work with domain experts

AI/ML Productization
End users

Make informed decisions
Find areas for model refinement

Public stakeholders

Voice & construct guidance for responsible use of AI

How AI is often biased.

Implicit bias

- A person with the bias does not aware of discrimination or prejudice against a person or group – whether it be on grounds of gender, race, disability, sexuality or class.

Sampling bias

- Randomly selected data does not reflect the distribution of the population. It is skewed towards some subset of the population.

Temporal bias

- A machine learning model works well currently but fails in the future. It didn't factor in possible future changes.

Over-fitting to training data

- An AI model adheres too much to the training dataset but cannot predict new data accurately. It does not generalize to a larger population.

Edge cases and outliers

- Data points are outside the normal distribution. Errors and noise are classified as edge cases: Errors are missing or incorrect values in the dataset; Normal variant

NOW

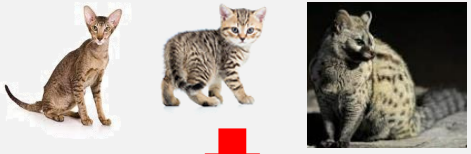

Training Data

Modeling & Training Process

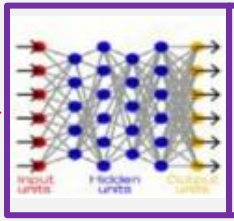

This is a cat
(p = .97)
(p = .80)

Learned Function

User

- Why did the model do that?
- Why not something else?
- When does the model succeed or fail?
- How do I correct an error?
- When can I trust the model?

FUTURE

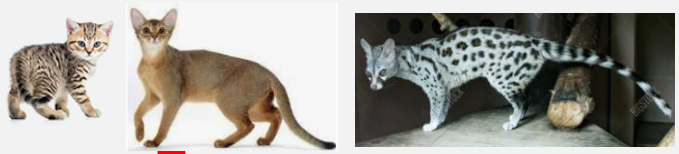

Training Data

New Modeling & Training Process

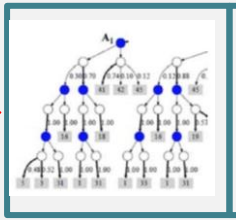

This is a cat because… :
- It has fur, whiskers, claws,…
- Its face has features…

Explainable Model

Explanation Interface

User

- I understand why
- I understand why not
- I know when the model will succeed or fail
- I know when to trust the model
- I know why the model erred

Let’s look at some examples

When Walmart Predicted What Shoppers Would Buy Before a Hurricane

Hurricane Charley struck the coast in August of **2004**

Daniel Ganninger Nov 4, 2020 · 2 min read ★

The top seller was beer (no surprise there), and strawberry Pop-Tarts, which sold seven times greater than normal.

With about 100 million customers that graced the stores each week, the company sat on a treasure trove of information. Walmart was believed to have 460 terabytes of data on its databases in the corporate headquarters in Bentonville, Arkansas, in 2004, which at the time, was enough data to double the amount of data on the Internet.

Forbes

Apr 21, 2017, 07:09am EDT

How Walmart Uses Data To Drive Disaster Relief

Dan Woods Former Contributor ⓘ
Enterprise & Cloud

Follow

This article is more than 4 years old.

Usually, when disaster strikes, whether it's natural like a hurricane or fire, or a terrorist attack, it seems our best preparations fail. Take Hurricanes Sandy and Katrina, which left whole swathes of the US devastated and uninhabitable for far longer than expected. The earthquake in Haiti a few years ago is another prime example where it took resources far too long to reach those in need.

Written by
Bernard Marr

Bernard Marr is a world-renowned futurist, influencer and

Walmart: Big Data analytics at the world's biggest retailer

23 July 2021

With over 20,000 stores in 28 countries, Walmart is the largest retailer in the world. So it's fitting then that the company is in the process of building the world's largest private cloud, big enough to cope with 2.5 petabytes of data every hour. To make sense of all this information, Walmart has created what it calls its Data Café – a state-of-the-art analytics hub located within its Bentonville, Arkansas headquarters.

49

Forbes

Feb 16, 2012, 11:02am EST

How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did

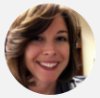

Kashmir Hill Former Staff
Tech
Welcome to The Not-So Private Parts where technology & privacy collide

Follow

This article is more than 9 years old.

Every time you go shopping, you share intimate details about your consumption patterns with retailers. And many of those retailers are studying those details to figure out what you like, what you need, and which coupons are most likely to make you happy. Target , for example, has figured out how to data-mine its way into your womb,

TARGET

"Target assigns every customer a Guest ID number, tied to their credit card, name, or email address that becomes a bucket that stores a history of everything they've bought and any demographic information Target has collected from them or bought from other sources."

"...women on the baby registry were buying larger quantities of unscented lotion around the beginning of their second trimester."

"Take a fictional Target shopper named Jenny Ward, who is 23, lives in Atlanta and in March bought cocoa-butter lotion, a purse large enough to double as a diaper bag, zinc and magnesium supplements and a bright blue rug. There's, say, an 87 percent chance that she's pregnant and that her delivery date is sometime in late August."

NEWS A timeline of Facebook's privacy issues – and its respo SHARE THIS –

facebook

Track buying habit of users over 40 websites

When: December 2007

What: Beacon, Facebook's first big brush with advertising privacy issues

Facebook's response: Zuckerberg apologizes, gives users choice to opt out

There was once a time when companies could track purchases by Facebook users and then notify their Facebook friends of what had been bought -- many times without any user consent.

When: July 2014

What: Mood-manipulation experiment on thousands of Facebook users

Facebook's response: Facebook data scientist apologizes

Facebook's mood-manipulation experiment in 2014 included more than half a million randomly selected users. Facebook altered their news feeds to show more positive or negative posts. The purpose of the study was to show how emotions could spread on social media. The results were published in the Proceedings of the National Academy of Sciences, kicking off a firestorm of backlash over whether the study was ethical.

AI, ML & DATA ENGINEERING

Facebook Announces ZionEX Platform for Training AI Models with 12 Trillion Parameters

Accelerating Facebook's infrastructure with application-specific hardware

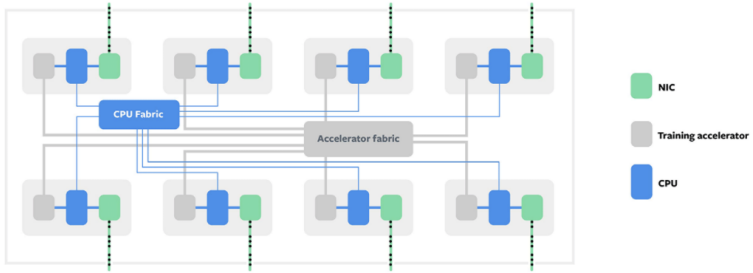

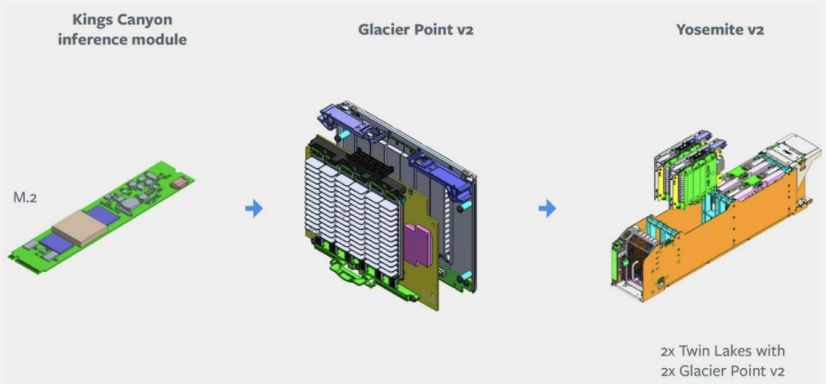

Application-specific hardware accelerators - Facebook Engineering (fb.com)

Can robots make up for Japan's care home shortfall?

Without more immigration, the country is forced to turn to technology

Grin and bear it: the Robear robot can lift elderly people or hospital patients off their bed into a wheelchair

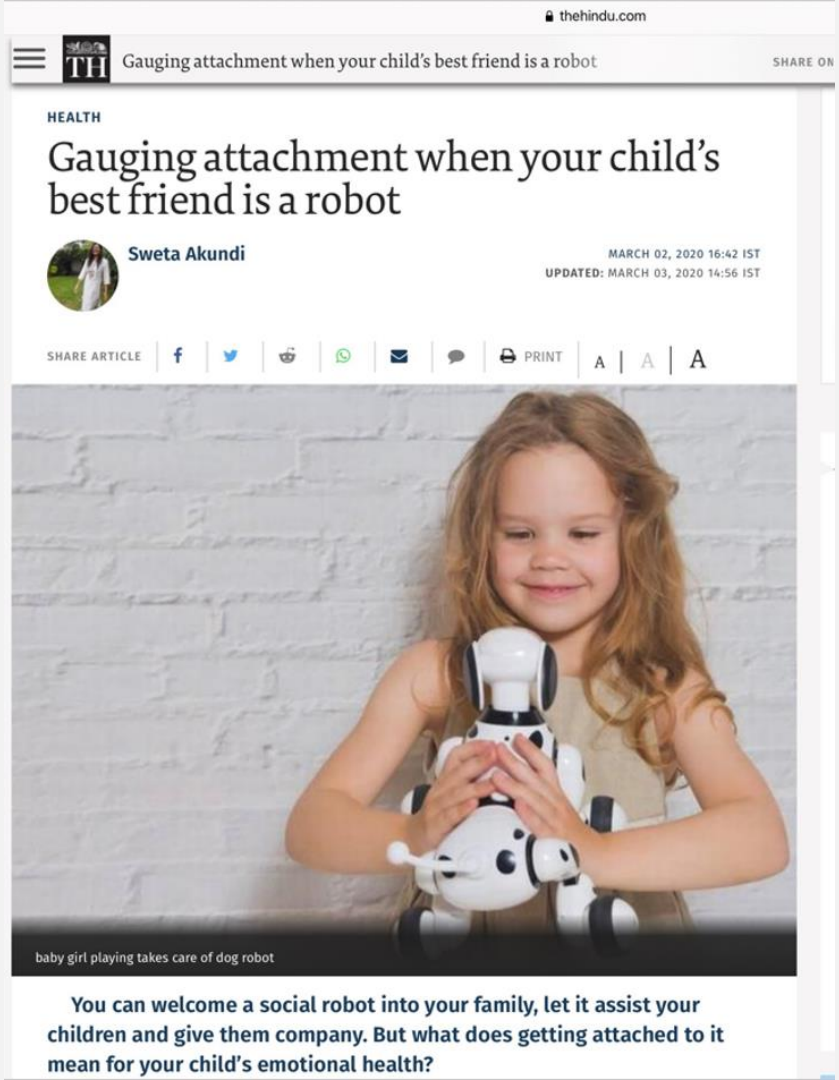

Can a child get too attached to a robot?

chatbotslife.com

ARTIFICIAL INTELLIGENCE

Will AI cause mass unemployment?

Are chatbots and self-driving cars going to make people lose their jobs?

alphaHoo Follow
May 29, 2019 · 6 min read

Artificial intelligence-based systems and chatbots are taking every industry by storm. The concern of AI taking over everyone's jobs is becoming increasingly urgent as recent AI breakthroughs attract public attention. As much as they are playing an important role in making tasks and processes better, their effect on human-centered jobs and capabilities in the

APRIL 6, 2016

Are robots taking our jobs?

by Moshe Y. Vardi, Rice University, The Conversation

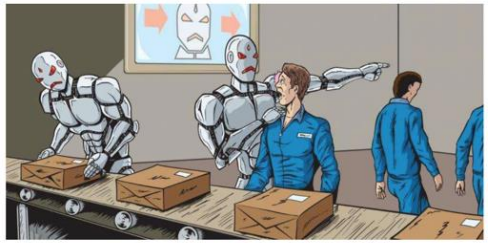

Is this a vision of the future? Credit: Robot worker image via shutterstock.com

If you put water on the stove and heat it up, it will at first just get hotter and hotter. You may then conclude that heating water

Featured Last Comments Popular

Implementing a 46-node quantum metropolitan area network
7 HOURS AGO 0

Hidden chamber found in Vanguard Cave – part of Gorham's Cave Complex in Gibraltar
7 HOURS AGO 0

orge.medium.com

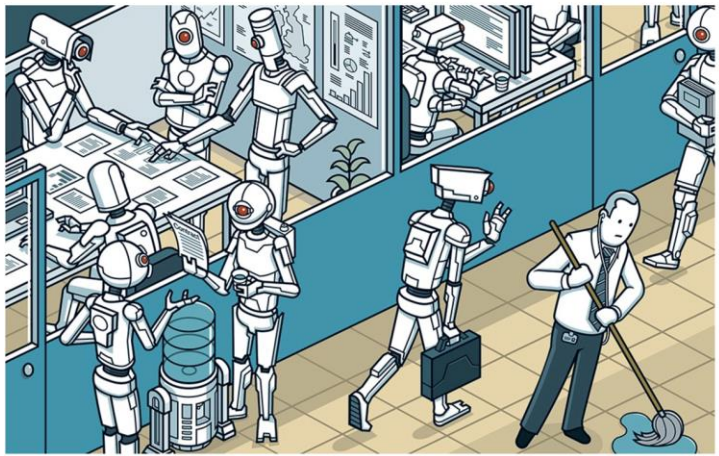

Illustration courtesy of: Nathan Daniels

The Future of Work: Technology Will Kill Your Job. Here's How

Morning Mix

Robot grabs man, kills him in German car factory

A VW Golf is assembled at the Volkswagen plant in Wolfsburg, northern Germany, in 2001. (Joerg Sarbach/AP)

By Abby Phillip

July 2, 2015

Some terrifying news out of Germany: A robot grabbed a worker and crushed him to death.

Who or what is the agent of responsibility?

Example : Mortgage Application

- Loan for a House
- Racial or class discrimination? (zip code)
- Algorithms used explainable?
  - ✓ Decision trees
  - ✓ Bayesian networks
  - ✓ Genetic algorithm
  - ✓ Neural networks & deep learning

Example : Medical Imaging & AI

- Imaging Modality Raw Data Set (MRI)

GBM (glioblastoma multiforme)

RCT

- Imaging as biomarker
- Pathology
- Radiology
- Genomics
- Precision medicine
- Targeted drug therapy

Traditional Approach

- Surgery
- Radiation therapy

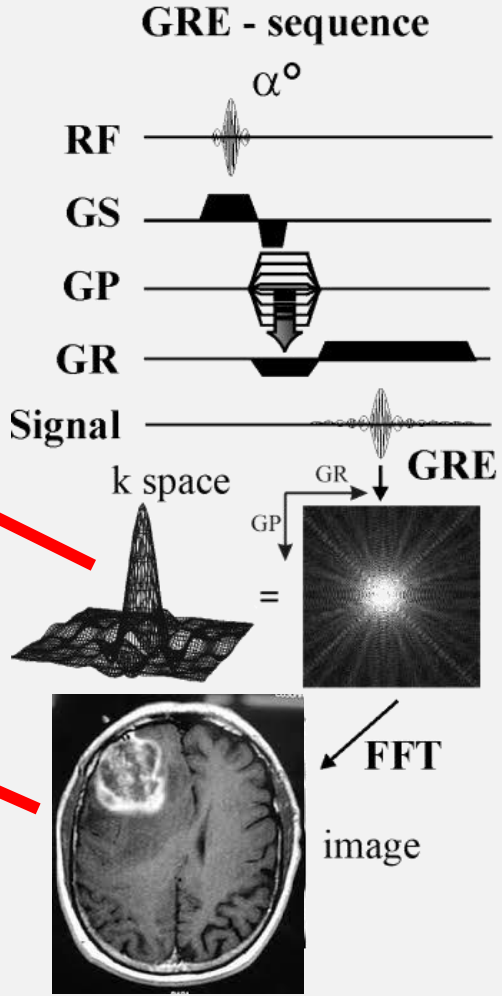

Source : Physics13_14.pdf (clinical-mri.com)

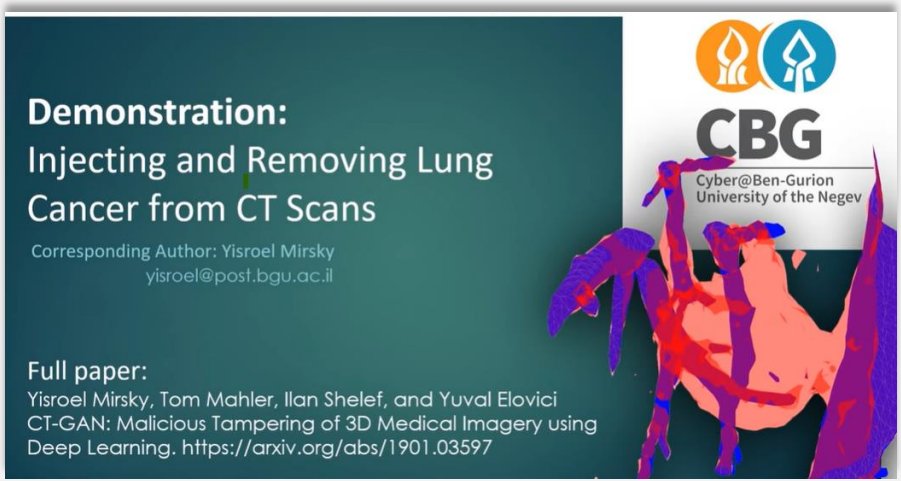

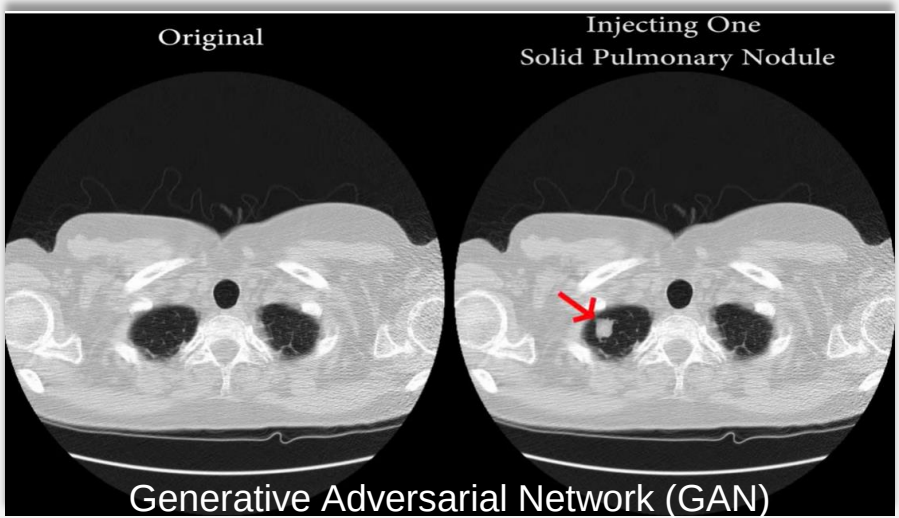

Cybersecurity

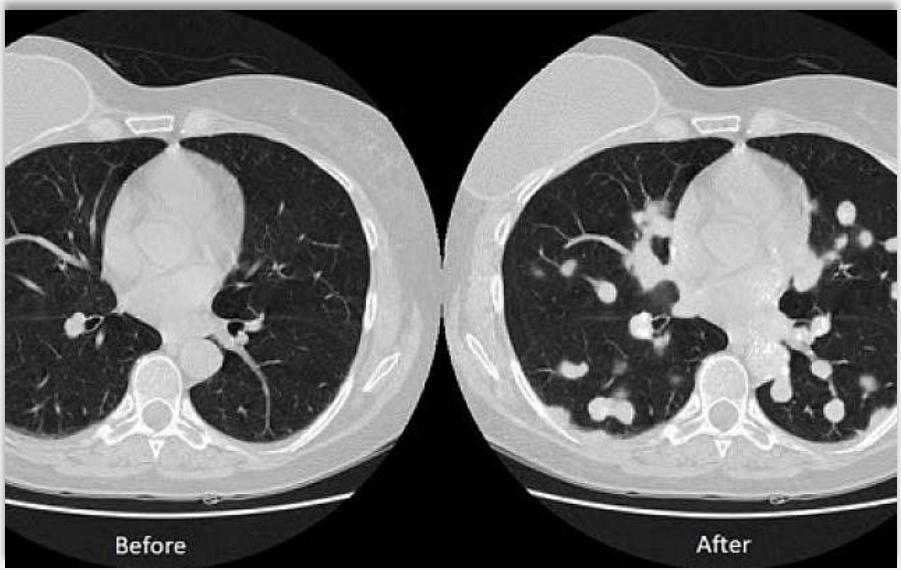

AI Detections

Wrongfully Accused by an Algorithm

In what may be the first known case of its kind, a faulty facial recognition match led to a Michigan man's arrest for a crime he did not commit.

Source : New York Times

Image "Cloaking" for Personal Privacy

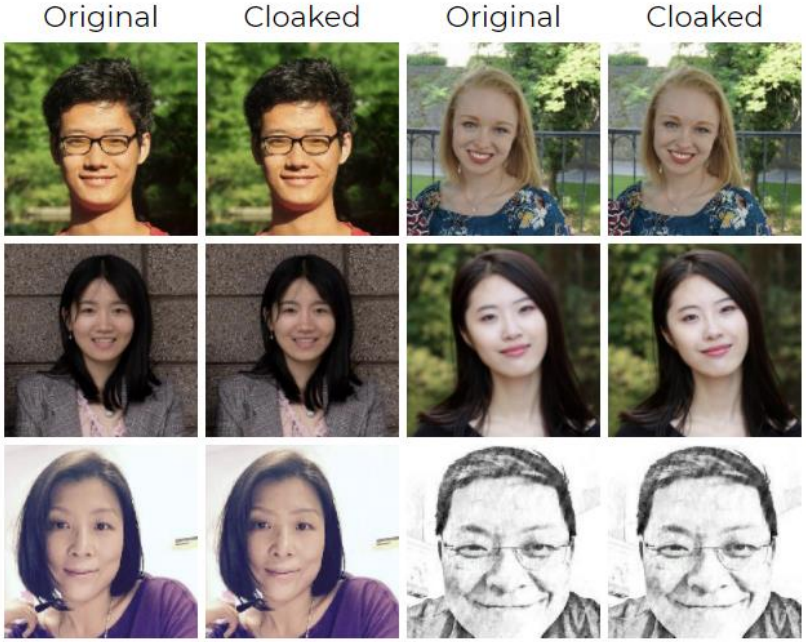

Shawn Shan†, PhD Student
Emily Wenger†, PhD Student
Jiayun Zhang, Visiting Student
Huiying Li, PhD Student
Haitao Zheng, Professor
Ben Y. Zhao, Professor
† *Project co-leaders and co-first authors*

- Email the Fawkes team
- Email us to join Fawkes mailing list for news on updates/changes.

SAND Lab
security, algorithms, networks and data

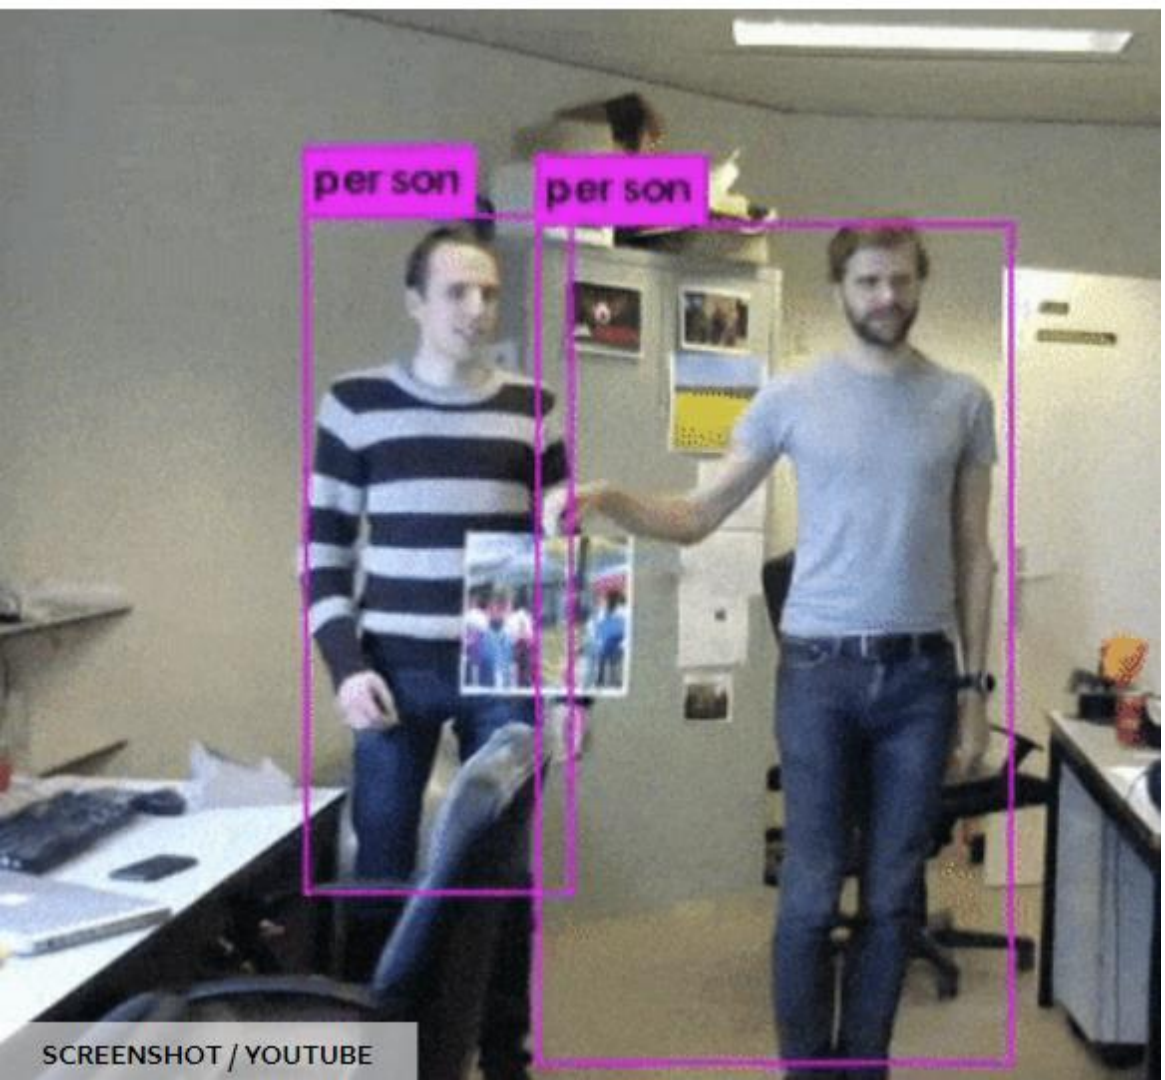

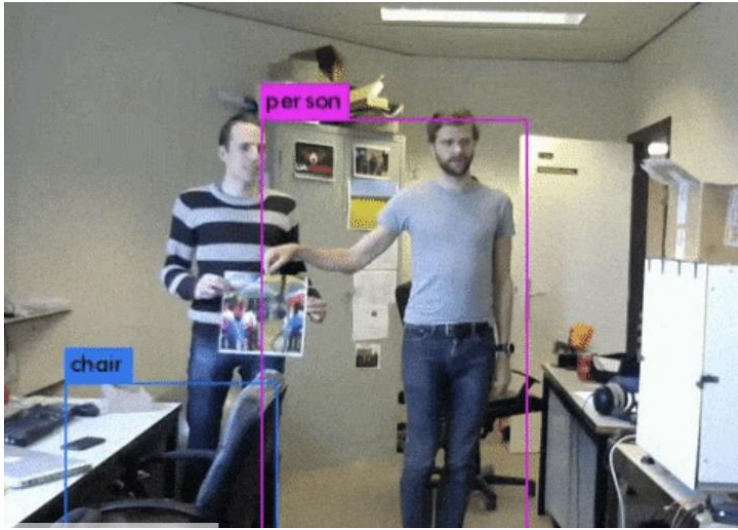

Source: Popular Mechanic

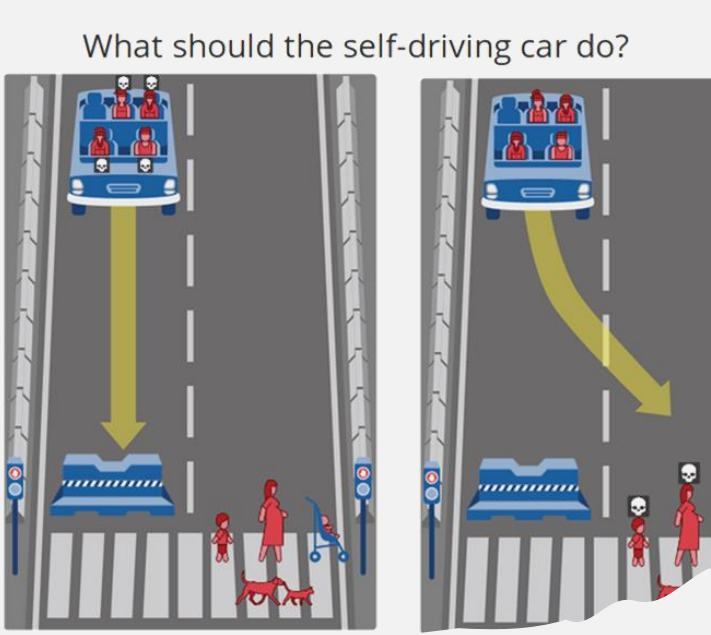

Autonomous Vehicles

Trolley problems and autonomous vehicles: what dos the public think? (bioedge.org)
Driverless cars and a new kind of “trolley problem” - Prospect Magazine
http://www.nexentire.com/webzine/201706/en/content_08.html

Ethics in AI : Qualities in Autonomous Driving

Technical and Non-Technical

- Human agency and oversight
- Technical robustness and safety
- Privacy and data governance
- Transparency
- Auditability
- Responsibility
- Predictability
- Incorruptibility (Cybersecurity & Hacking)
- Accountability

63


CNN Money
This is the first known pedestrian fatality involving a self-driving car

Tesla Autonomous Driving & Explainable AI

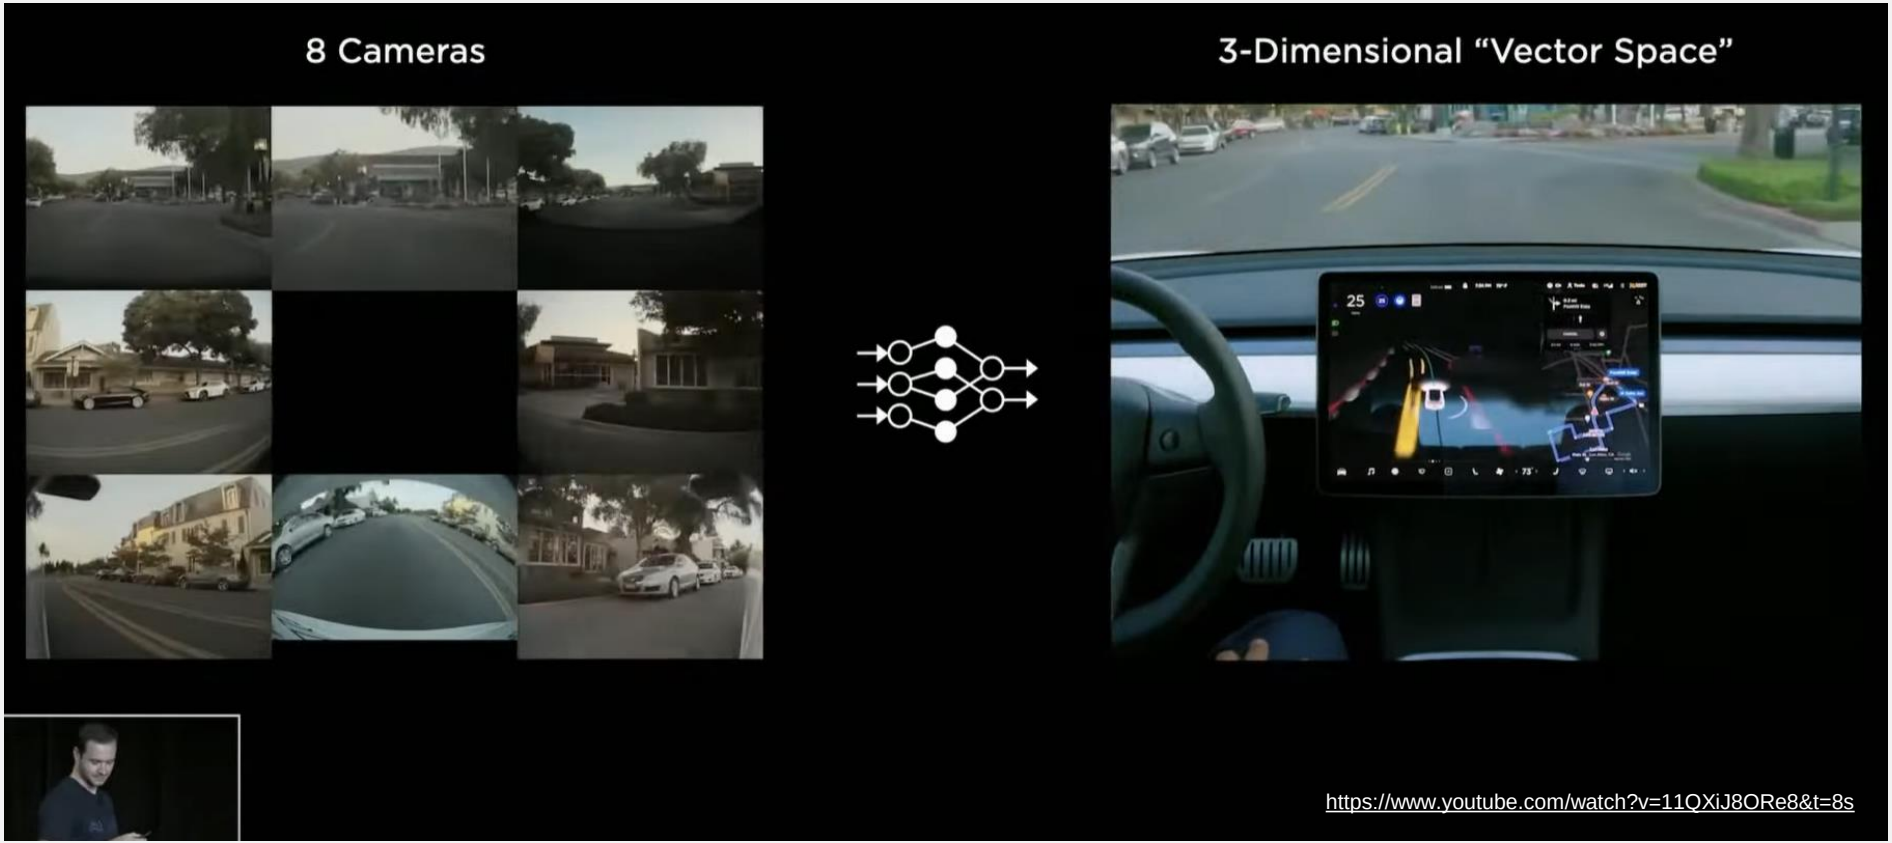

Ethical issues : Privacy & Anonymization & Safety

Model Creation

Biological Visual Cortex Wiring

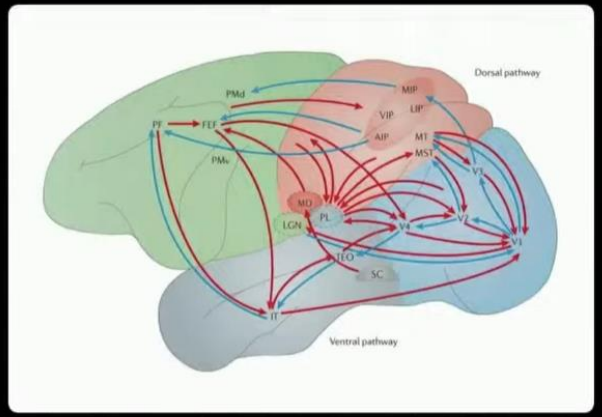

Top-down influences on visual processing. Nature Reviews Neuroscience. 2013.

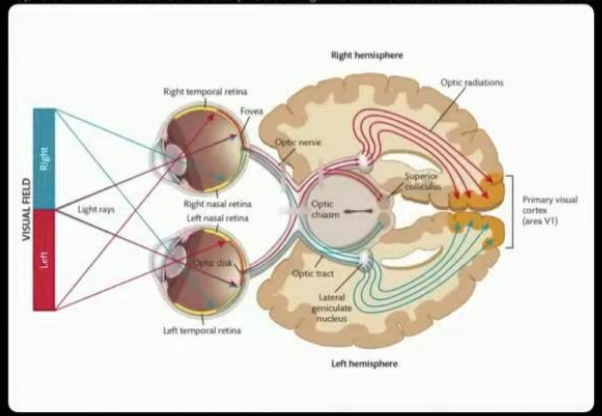

https://quizlet.com/ca/268598265/primary-visual-cortex-5th-diagram/

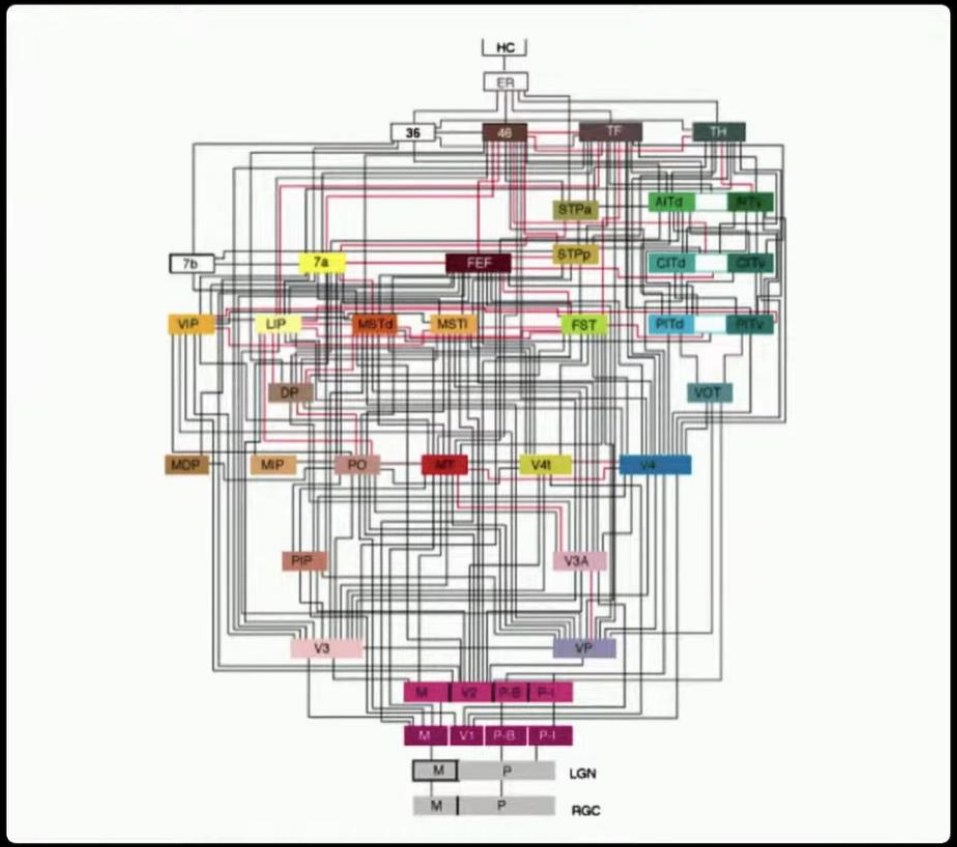

Felleman DJ, Van Essen DC. Distributed hierarchical processing in the primate cerebral cortex. Cereb Cortex. 1991 Jan-Feb;1(1):1-47.

Camera Input

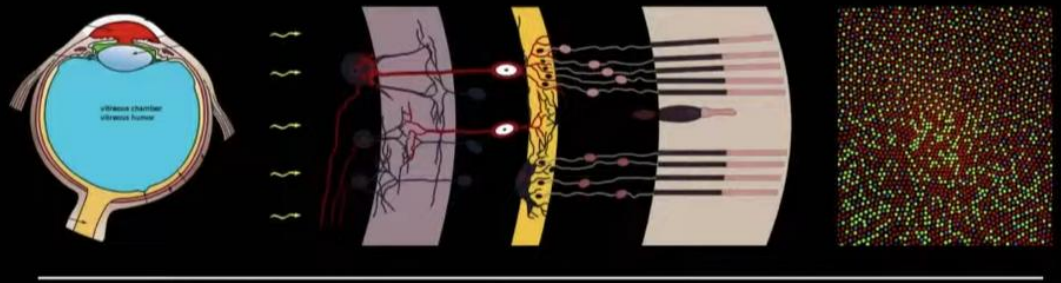

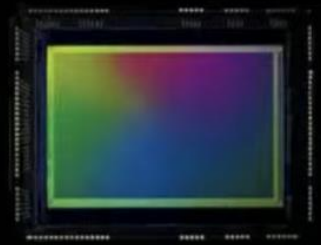

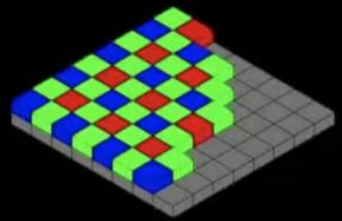

raw

1280x960 12-Bit (HDR) @ 36Hz

CCD

Pixels

- Estimated Autopilot miles to-date: **3.3 billion** miles
- Estimated miles in all Tesla vehicles: **22.5 billion** miles

Source: Tesla Vehicle Deliveries and Autopilot Mileage Statistics - Lex Fridman

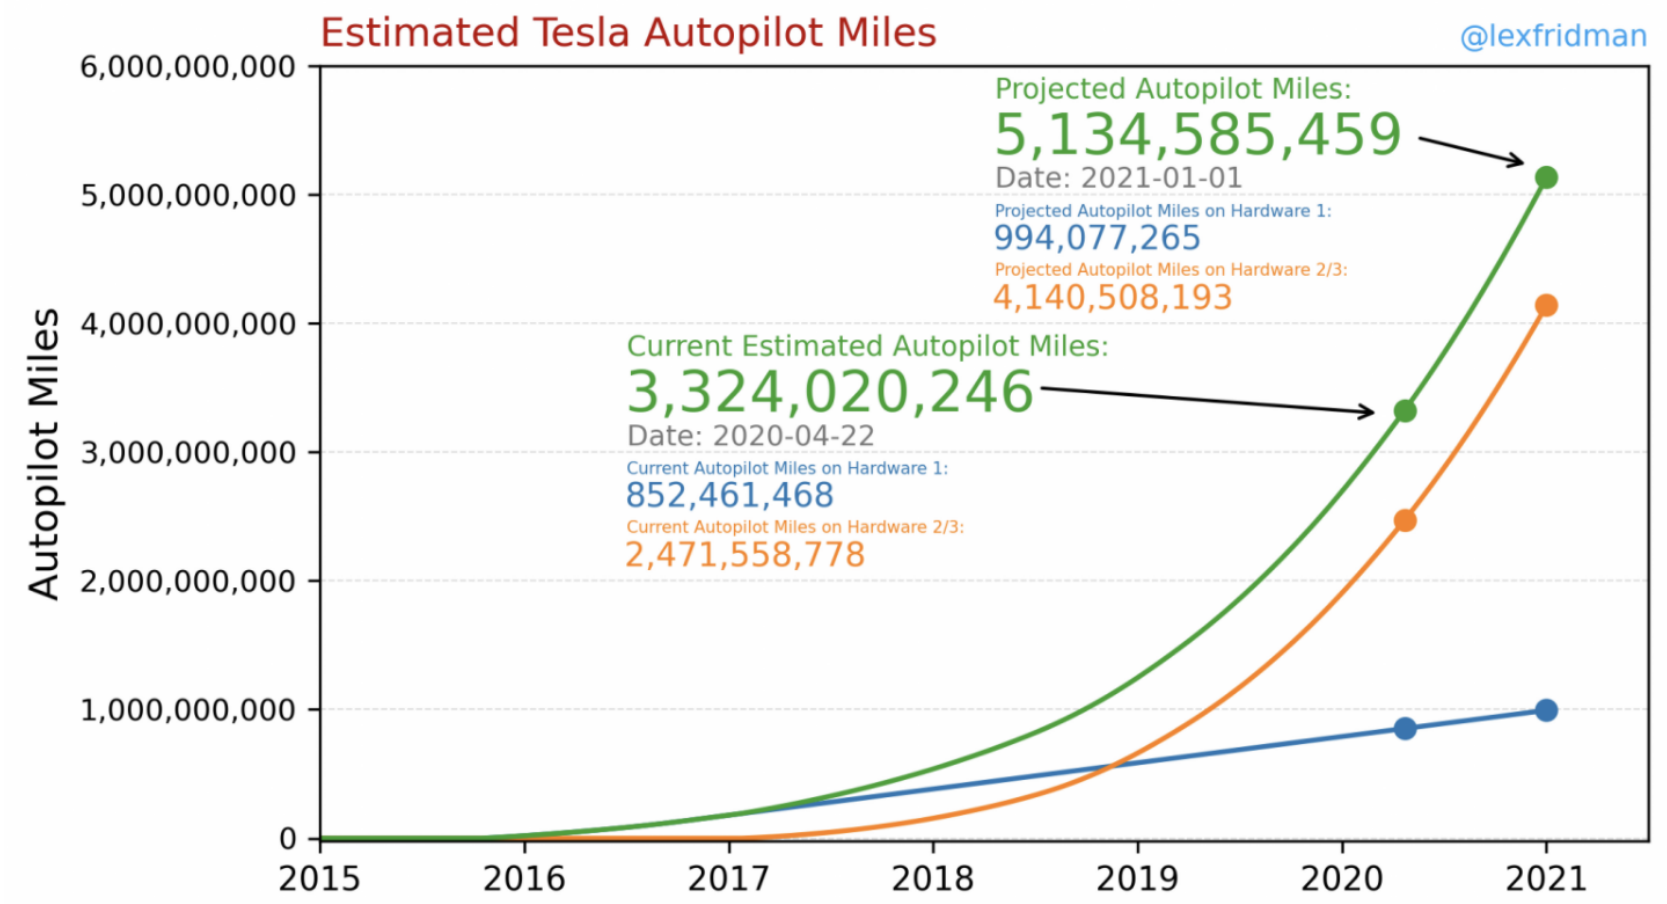

Neural Network Backbone

Residual Neural Network (RegNet)

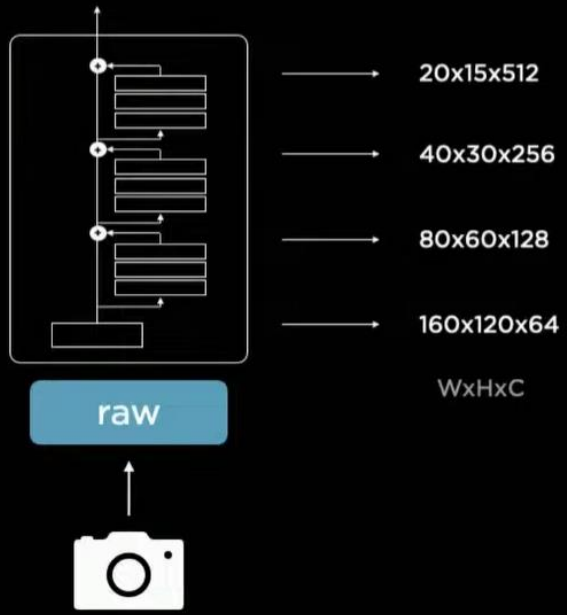

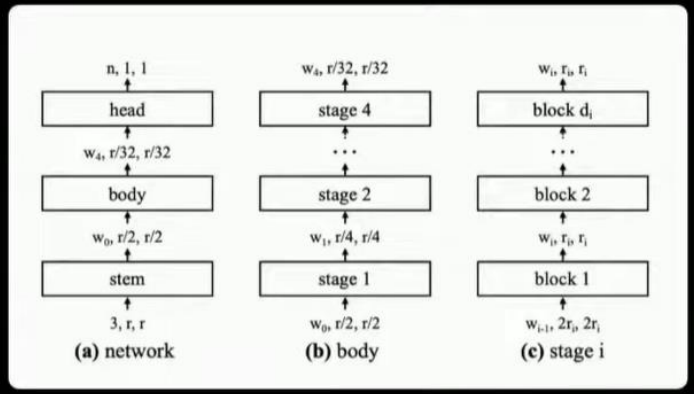

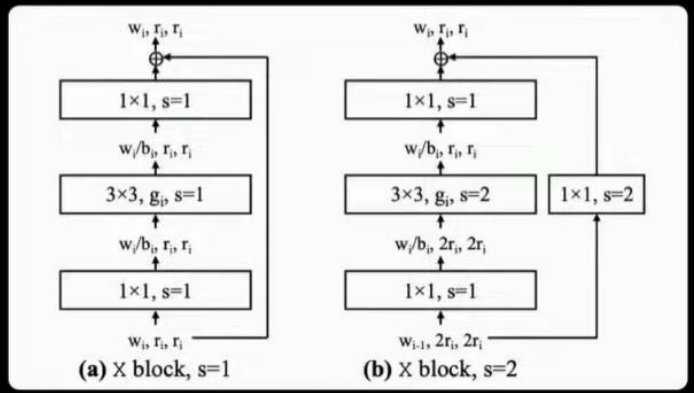

Designing Network Design Spaces, Radosavovic et al. 2020

Multi-Scale Feature Pyramid Fusion

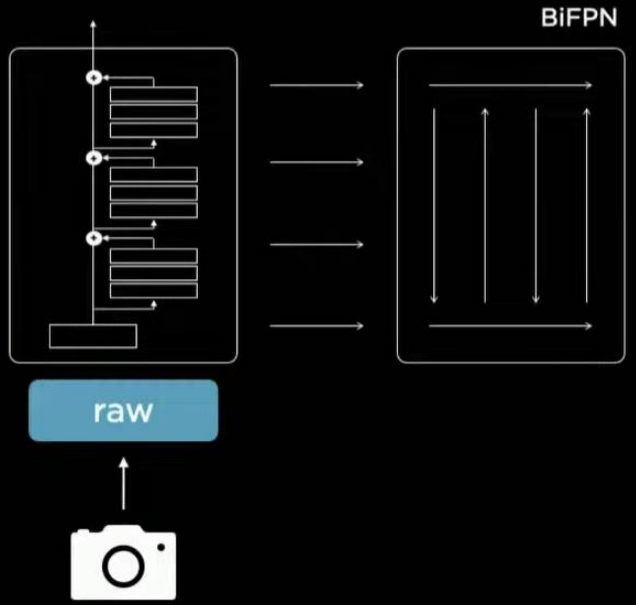

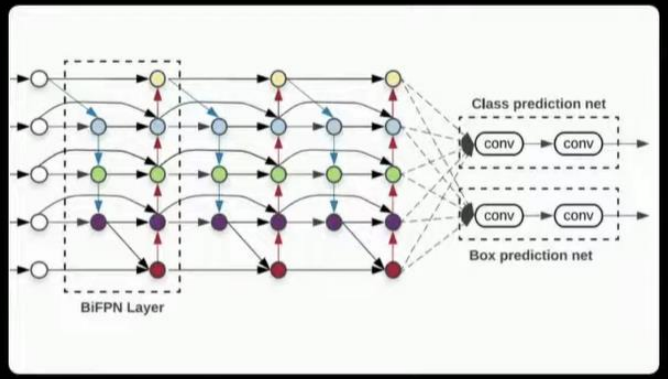

EfficientDet: Scalable & Efficient Object Detection, Tan et al. 2019

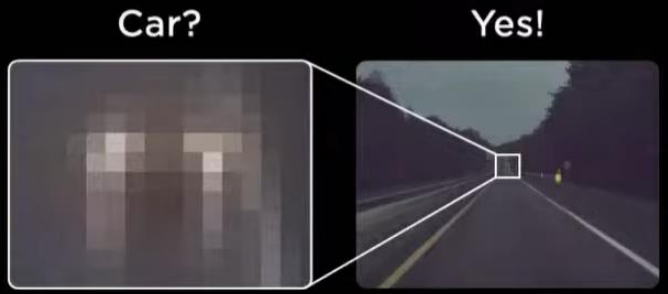

Detection Head

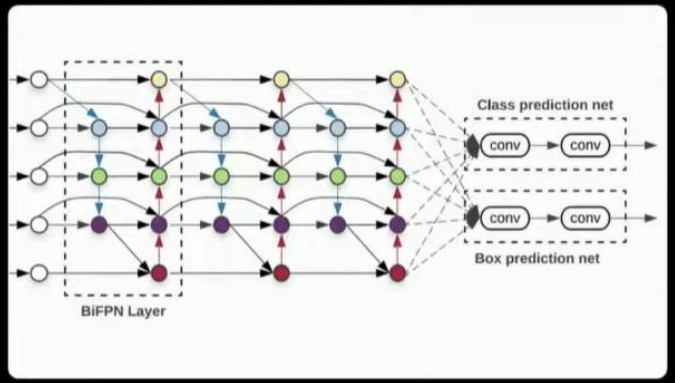

EfficientDet: Scalable & Efficient Object Detection, Tan et al. 2019

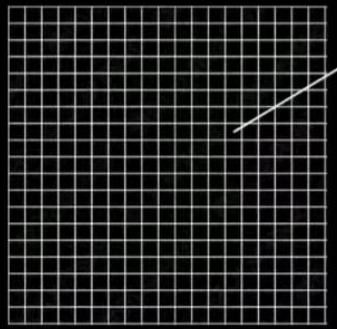

For Each Pixel:

- Is there an object here?
- If there were an object here, what is its extend and attributes?

Multi-Task Learning “HydraNets”

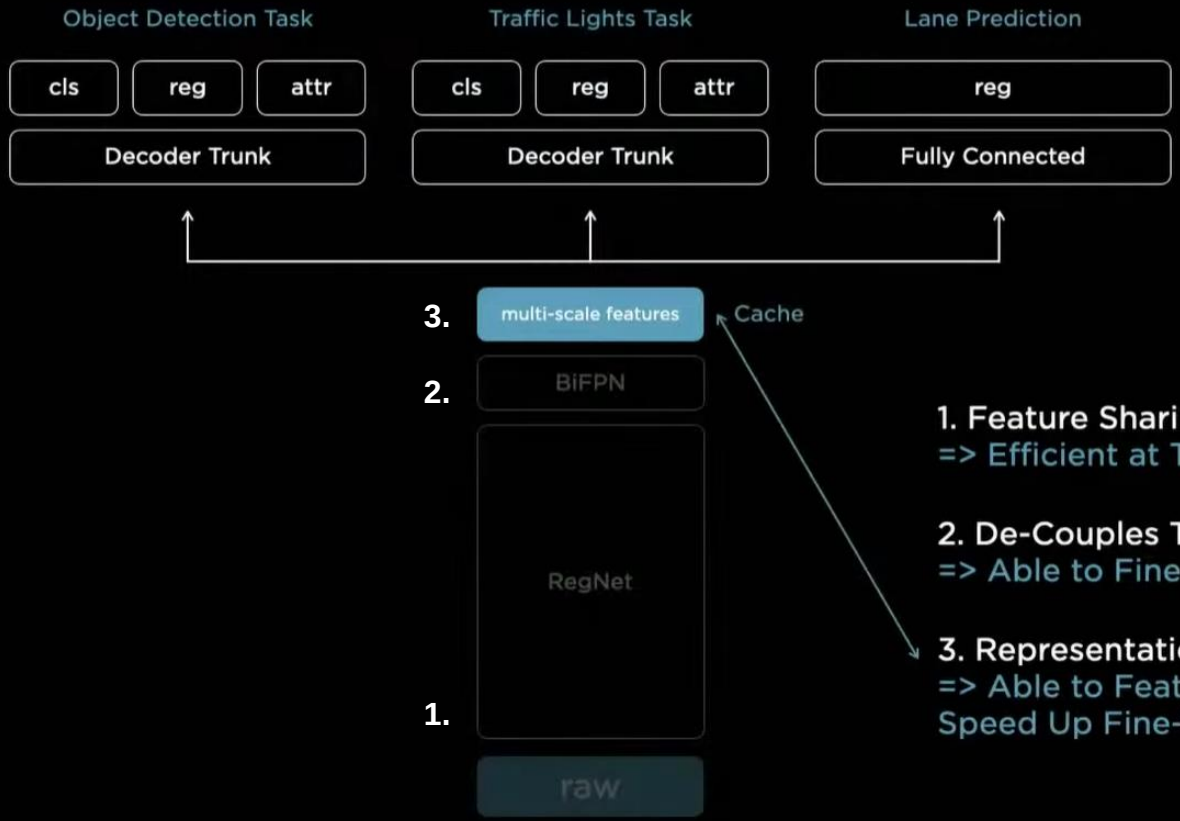

1. Feature Sharing
=> Efficient at Test Time

2. De-Couples Tasks
=> Able to Fine-Tune Tasks Individually

3. Representation Bottleneck
=> Able to Feature Cache &
Speed Up Fine-Tuning


Ego Speed: 45.56 MPH
time: 1545.441522000
CAL P 0.60 Y 1.20 R 0.00 deg
Vision fps: 18.05 Draw fps: 17.67 Display fps: 21.34
NL(0.00), E(0.93), F(0.07), TF(0.00), S(0.00)
NRW: FLP(0.00), FRP(0.00)
CutInExcited (Prb 0.56)
+0.0001 AUTO_HIGH_BEAM
+0.0000 BLINDED
+0.0002 RAINING
+0.0000 TIRE_SPRAY
+0.0013 WET_ROAD
0.7902 RESTRICTED
0.0934 CONTROLLED_ACCESS
AP
L:0 R:0 F:2 ON:0
W:8.2 AP:1.0 I:0
VS: 46.7 MPH St: 1
merge: 1.0 1 160.2 R
40m
MAIN –
TESLA LIVE

Why
Plan Jointly?

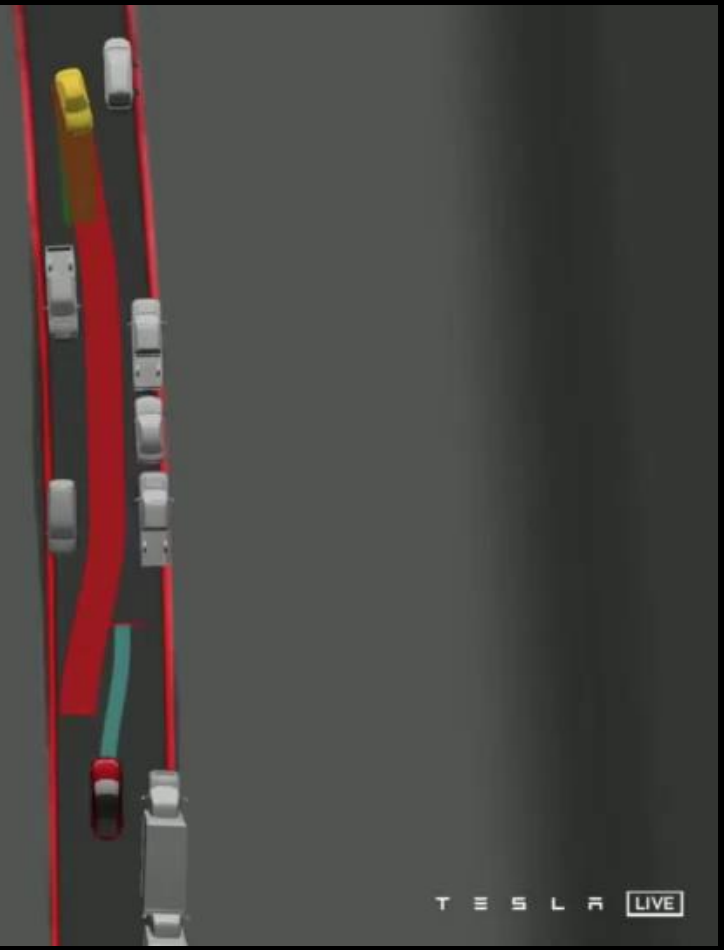

The Final Architecture

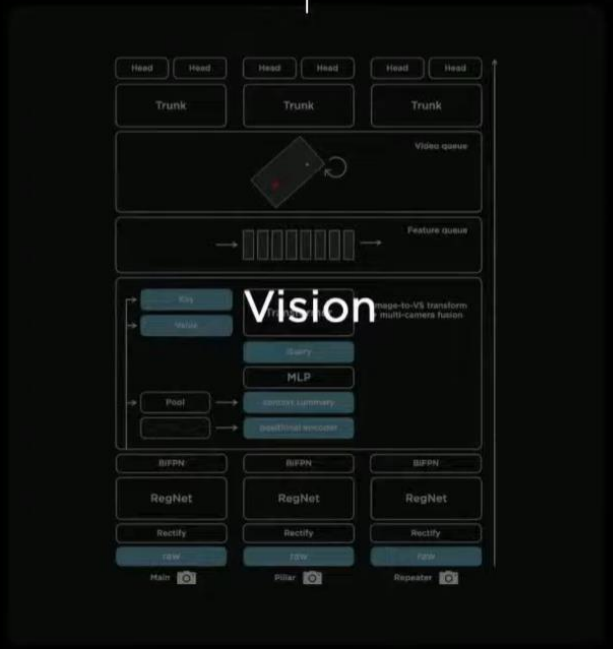

Vector Space

Intermediate Features

Explicit Planning & Control

Trajectory Distribution

Neural Net Planner

Steering & Accel

HOW DO WE MAKE A CAR AUTONOMOUS?

HOW DO WE GENERATE TRAINING DATA?

HOW DO WE RUN IT IN THE CAR?

HOW DO WE ITERATE QUICKLY?

Manual Labeling

Auto Labeling

Simulation

Scaling Data Generation

Life of a Clip

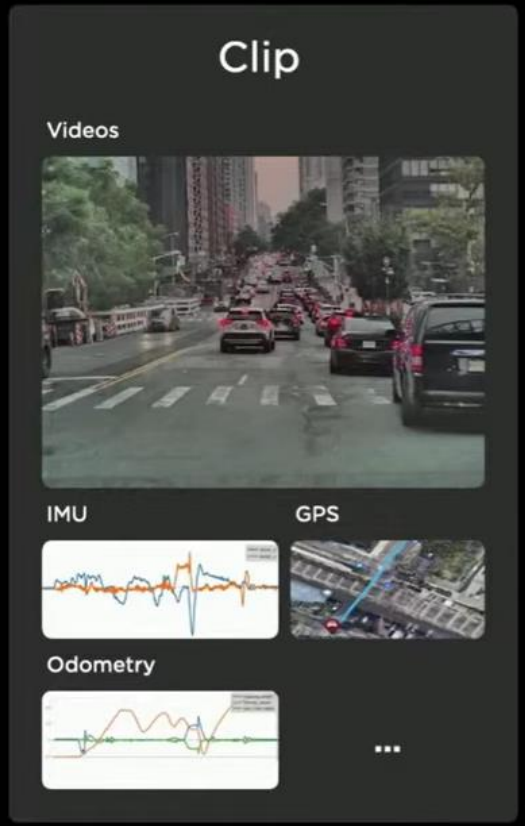

Offline Neural Networks

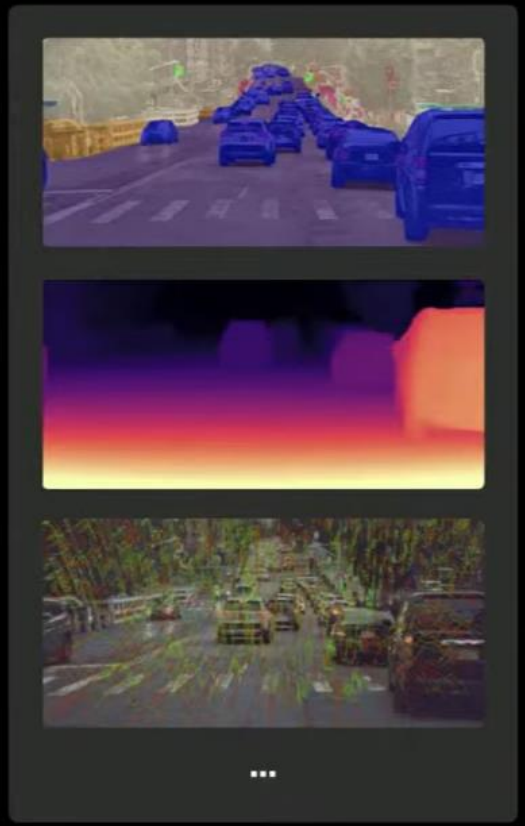

Ego Trajectory & Static World Reconstruction

Moving Objects & Kinematics

Data Labeling Growth

1,000 Person In-House Labeling Team

Fully Custom Built Data Labeling & Analytics Infrastructure

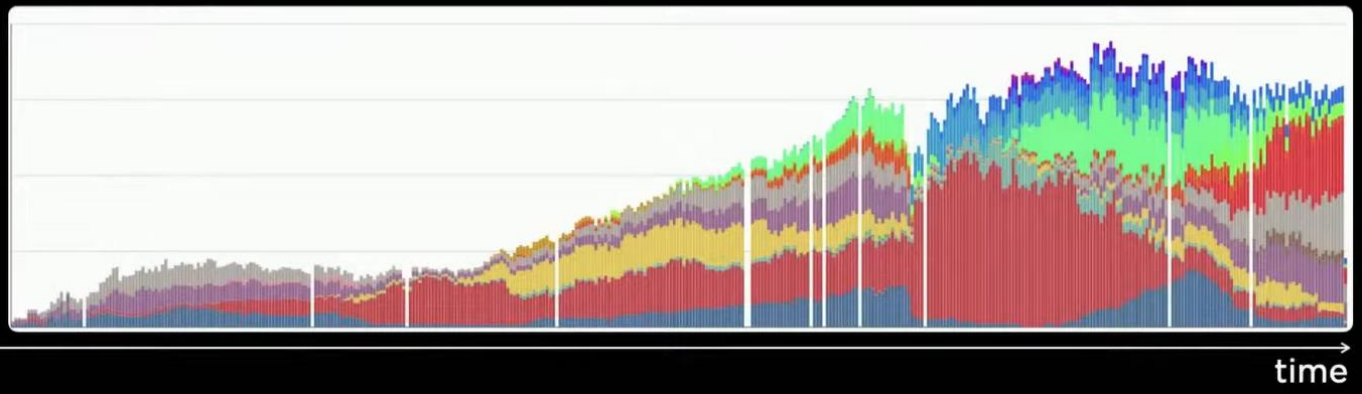

Simulation Is a Video Game With Autopilot as the Player

TESLA LIVE

Hardware Integration

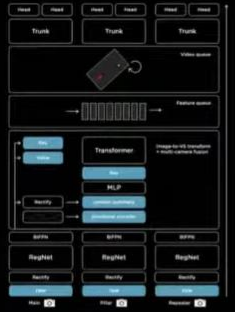

Neural Networks

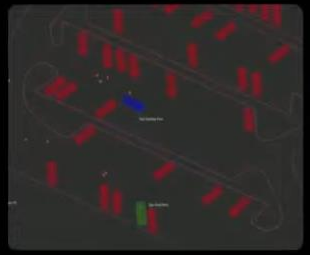

Algorithms

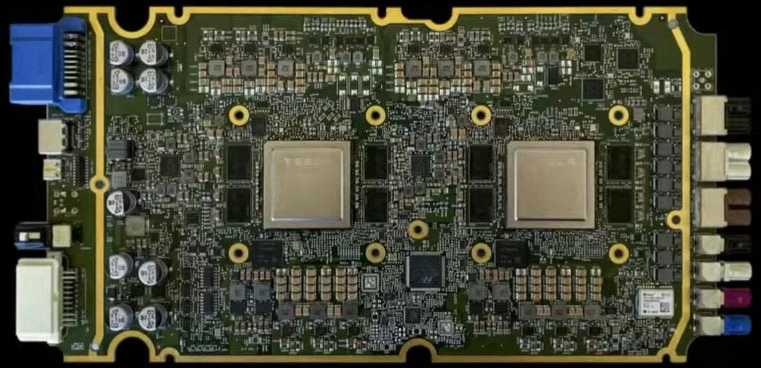

FSD Computer

Minimize Latency
&
Maximize Frame Rate

Training Node Architecture

Superscalar In-Order CPU
4 Wide Scalar + 2 wide Vector Pipes

4-Way Multithreaded

Custom ISA Optimized for ML Kernels

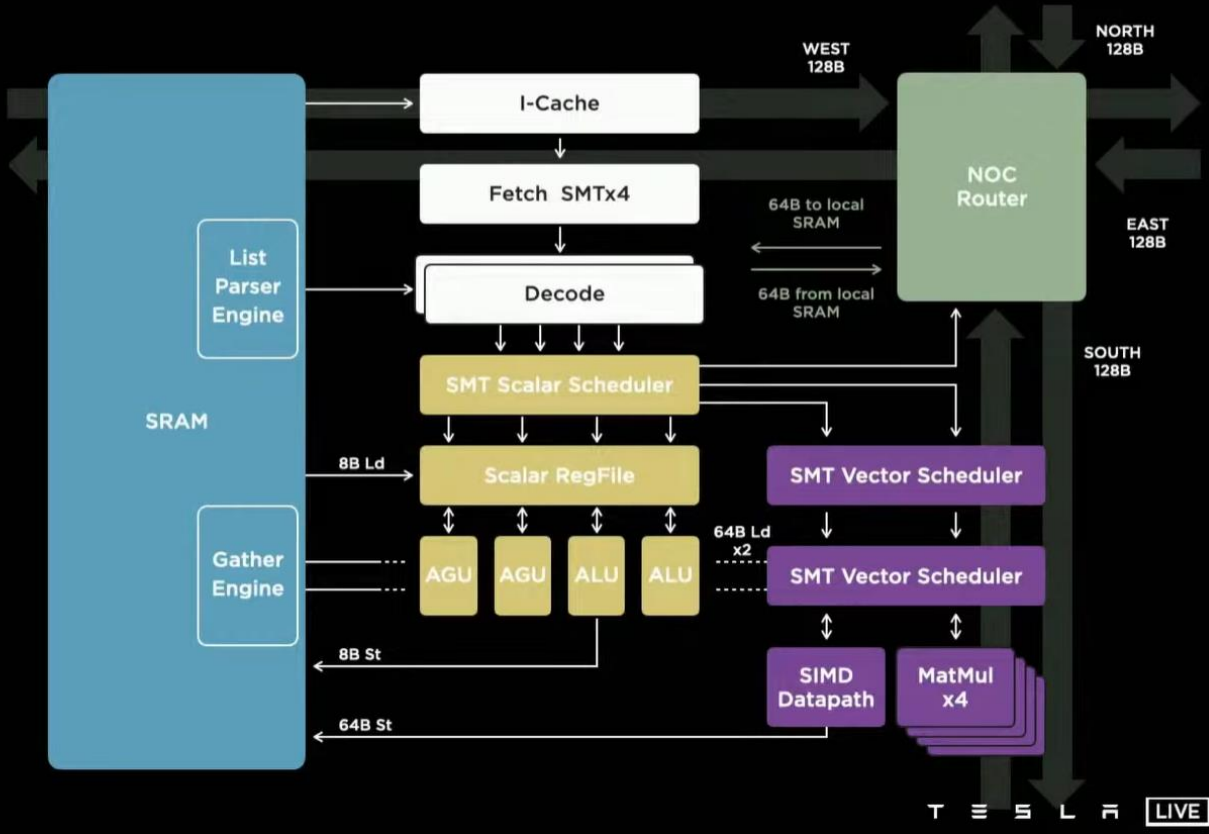

D1 Chip

362 TFLOPs BF16/CFP8
22.6 TFLOPs FP32

10TBps/dir. On-Chip Bandwidth
4TBps/edge. Off-Chip Bandwidth

400W TDP

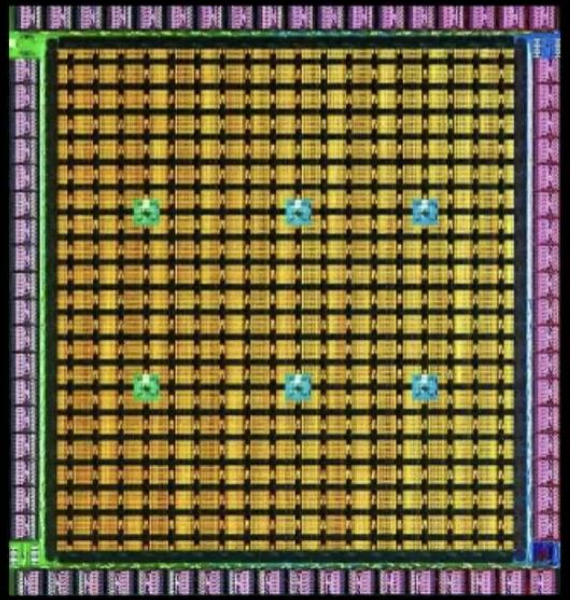

645mm²
7nm Technology

50 Billion
Transistors

11+ Miles
Of Wires

Training Tile

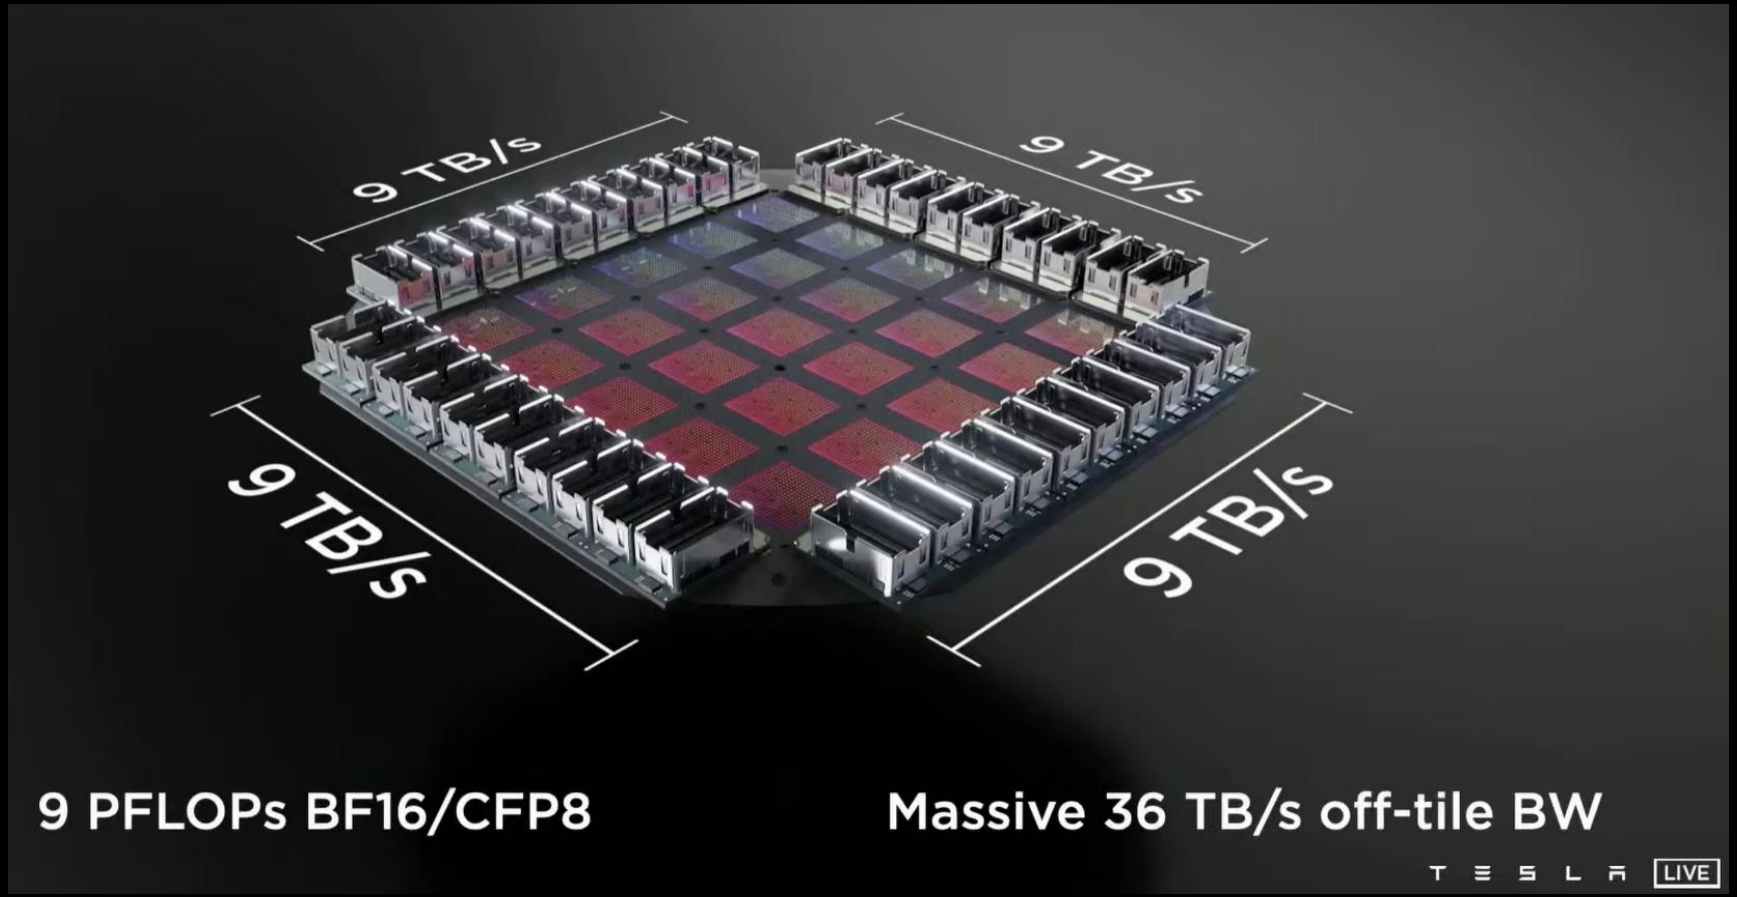

Training Tile

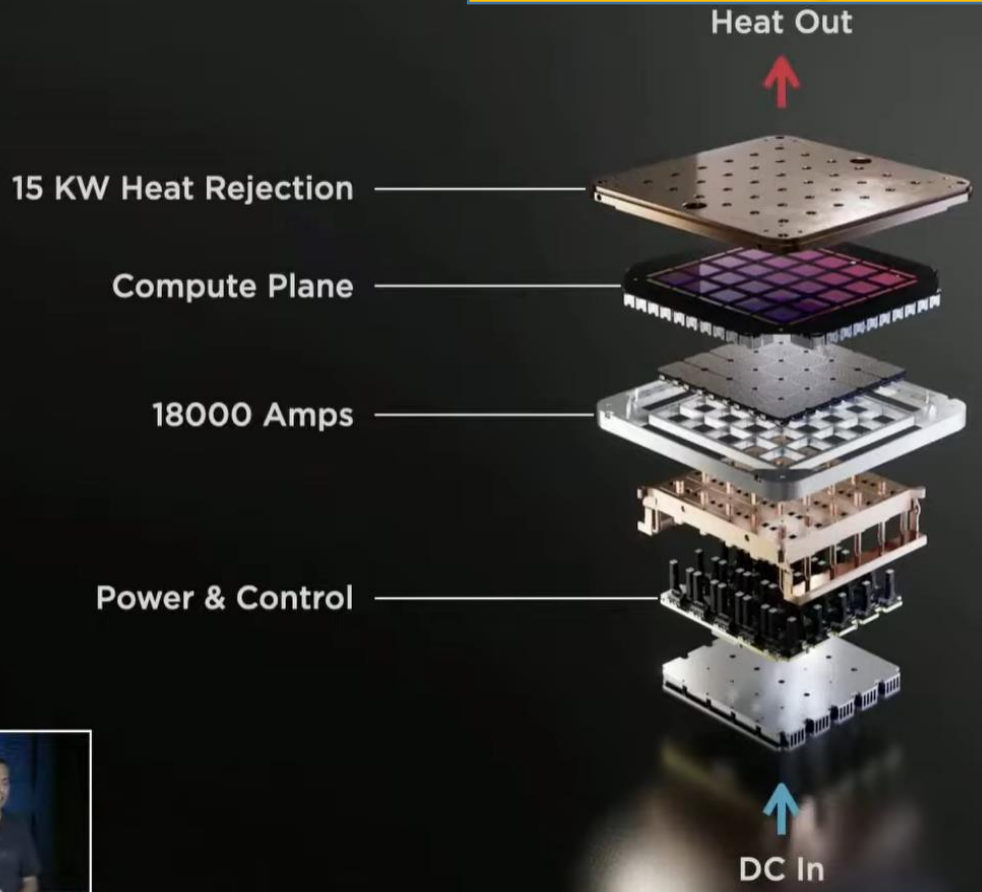

Performance & Scale at all Levels

- Software
- Compute Cluster
- System
- Chip

Logical View of the System

Distributed System Is Partitionable

DPU - Dojo Processing Unit

A Virtual Device That Can Be Sized According to Application Needs

D1 Accelerator Chips (Compute + Local Memory)
Dojo Interface Processors (Ingest + Shared Memory)

ExaPOD


1.1 EFLOP (BF16/CFP8)
120 TRAINING TILES | 3000 D1 CHIPS | >1M TRAINING NODES
Uniform High BW
& Low-Latency Fabric
TESLA LIVE

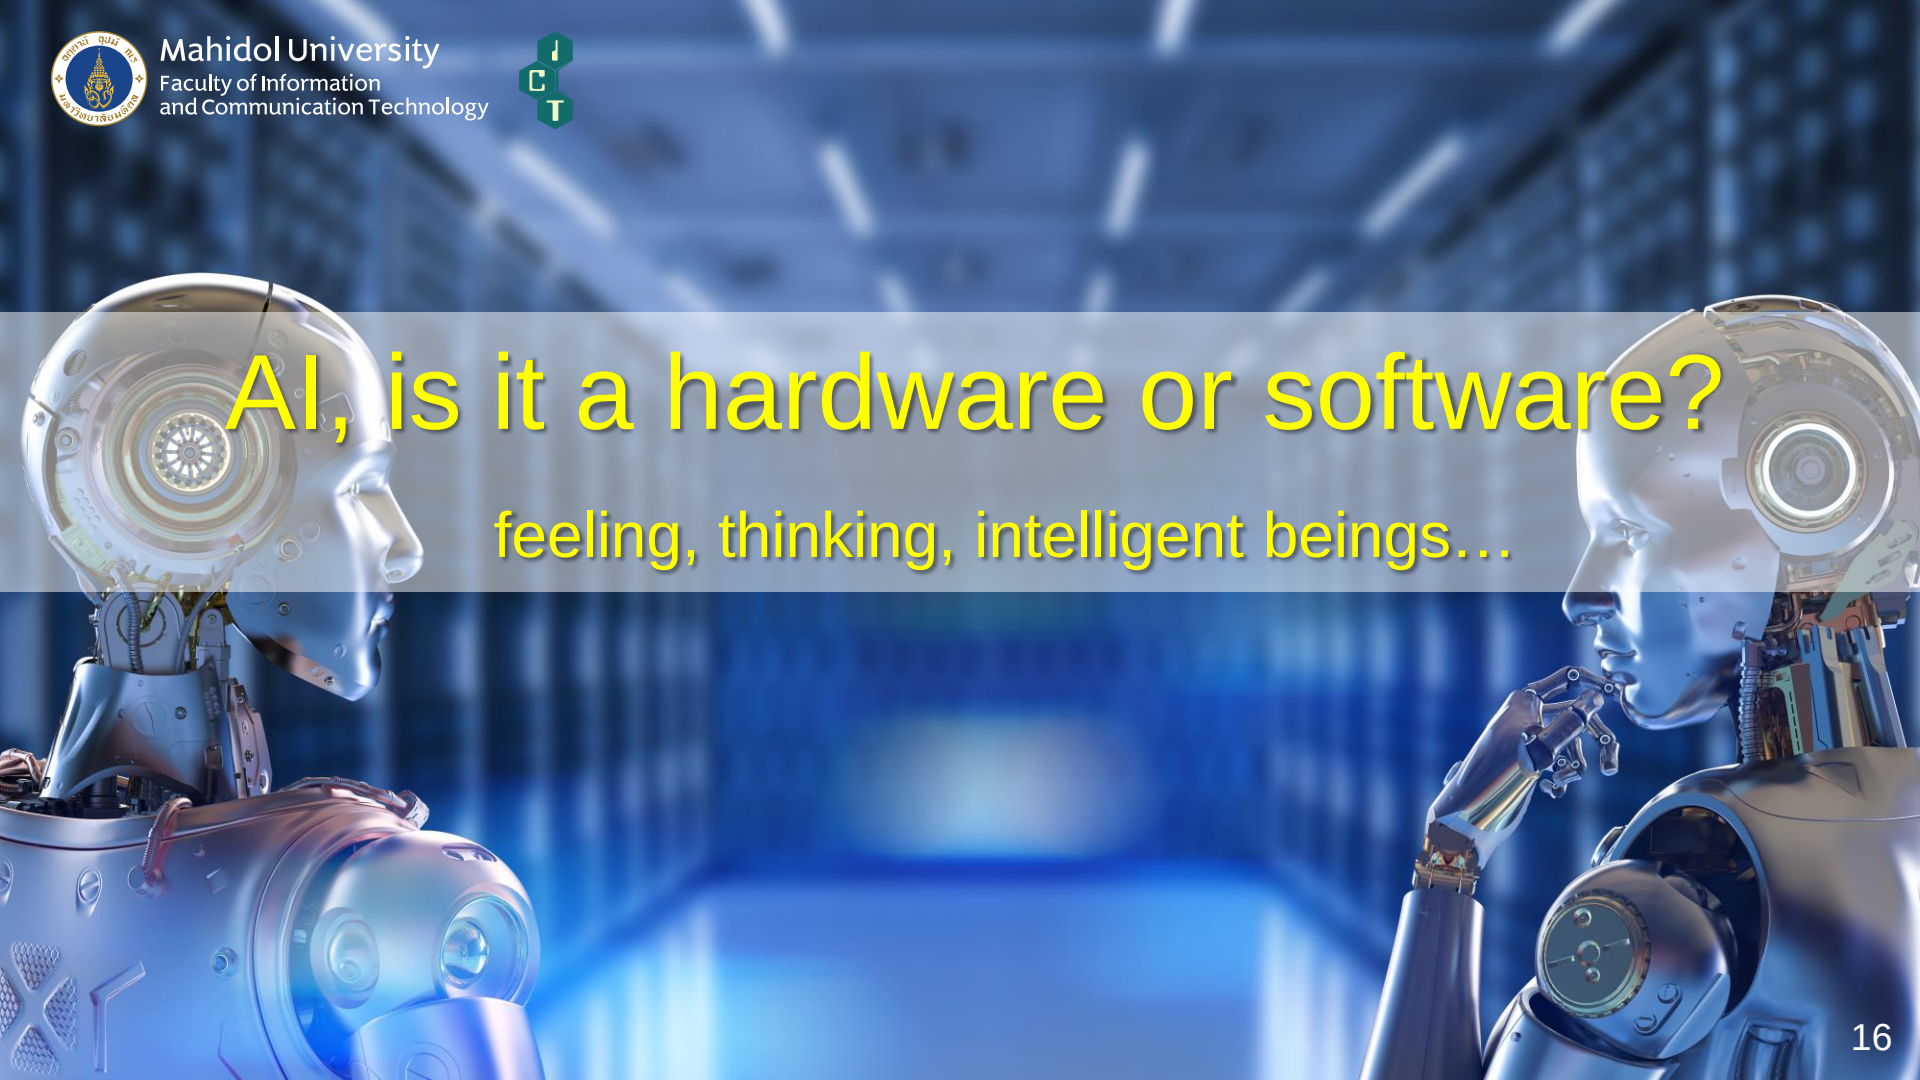

AI, is it a hardware or software?

feeling, thinking, intelligent beings…

Concluding Remarks

A.I. Research

- Code of Ethics
  - ✓**Explainability**
  - ✓**Auditability**
  - ✓**Transparency : disclose any actual or potential danger to the public**
  - ✓**Responsibility for their own work**
  - ✓**Use AI only if believed to be safe**
  - ✓**Fairness**
  - ✓**Security**
  - ✓**Gather and address grave public concern caused by AI**
  - ✓**Policy**
  - ✓**Privacy (PDPA, HIPPA, GDPR)**

No one size fit all solution.

Domain and content specific…

Where to training…

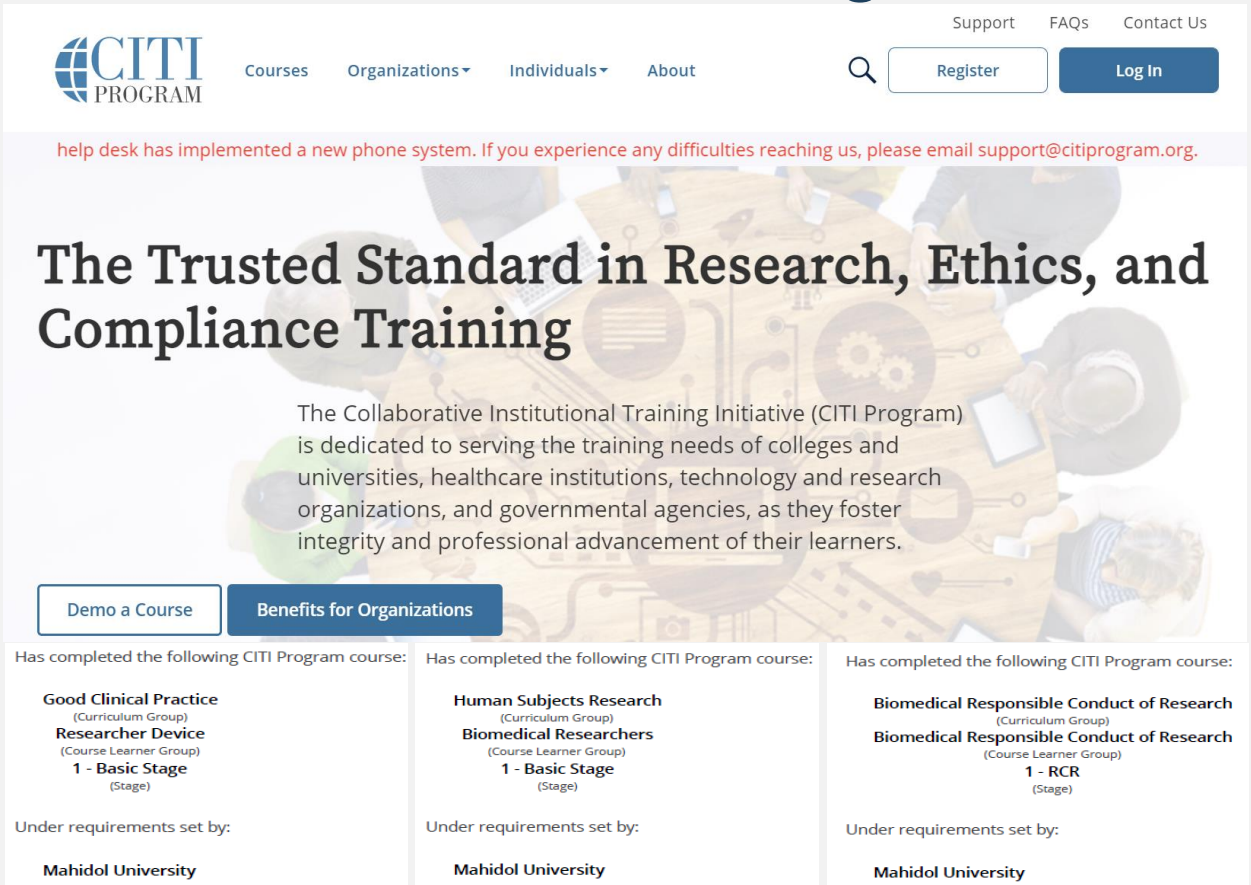

Research, Ethics, and Compliance Training | CITI Program

Incoming AI Resources

98

Mahidol AI Center

However, running AI experiments in medical imaging demand powerful hardware as well because using AI in medical imaging involves intensive tasks like training it to detect lesions and other signs of cancer in thousands of scans. Here is a difference:

Training AlexNet – A deep learning neural network for image recognition

Xeon E5 16 core CPU
US$3,500

43 days

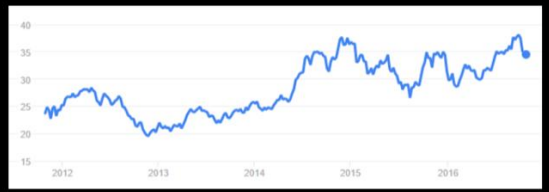

Titan X GPU
US$999

2.5 days

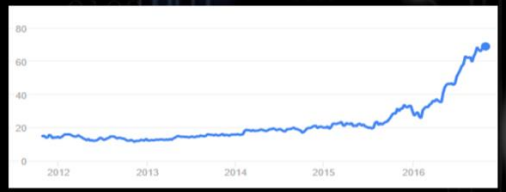

1. 8X NVIDIA A100 GPUS WITH UP TO 640 GB TOTAL GPU MEMORY
   12 NVLinks/GPU, 600 GB/s GPU-to-GPU Bi directonal Bandwidth

2. 6X NVIDIA NVSWITCHES
   4.8 TB/s Bi-directional Bandwidth, 2X More than Previous Generation NVSwitch

3. 10X MELLANOX CONNECTX-6 200 Gb/s NETWORK INTERFACE
   500 GB/s Peak Bi-directional Bandwidth

4. DUAL 64-CORE AMD CPUs AND 2 TB SYSTEM MEMORY
   3.28 More Cores to Power the Most Intensive Al Jobs

5. 30 TB GEN4 NVME SSD|
   50 GB/s Peak Bandwidth, 2x Faster than Gen3 NVME SSDS

	A100 40GB PCIe	A100 80GB PCIe	A100 40GB SXM	A100 80GB SXM
FP64	9.7 TFLOPS			
FP64 Tensor Core	19.5 TFLOPS			
FP32	19.5 TFLOPS			
Tensor Float 32 (TF32)	156 TFLOPS | 312 TFLOPS*			
BFLOAT16 Tensor Core	312 TFLOPS | 624 TFLOPS*			
FP16 Tensor Core	312 TFLOPS | 624 TFLOPS*			
INT8 Tensor Core	624 TOPS | 1248 TOPS*			
GPU Memory	40GB HBM2	80GB HBM2e	40GB HBM2	80GB HBM2e
GPU Memory Bandwidth	1,555GB/s	1,935GB/s	1,555GB/s	2,039GB/s
Max Thermal Design Power (TDP)	250W	300W	400W	400W
Multi-Instance GPU	Up to 7 MIGs @ 5GB	Up to 7 MIGs @ 10GB	Up to 7 MIGs @ 5GB	Up to 7 MIGs @ 10GB
Form Factor	PCIe		SXM	
Interconnect	NVIDIA® NVLink® Bridge for 2 GPUs: 600GB/s ** PCIe Gen4: 64GB/s		NVLink: 600GB/s PCIe Gen4: 64GB/s	
Server Options	Partner and NVIDIA-Certified Systems™ with 1-8 GPUs		NVIDIA HGX™ A100-Partner and NVIDIA-Certified Systems with 4,8, or 16 GPUs NVIDIA DGX™ A100 with 8 GPUs	

Mahidol AI Center Empowers Advanced Innovations & Future Expansion

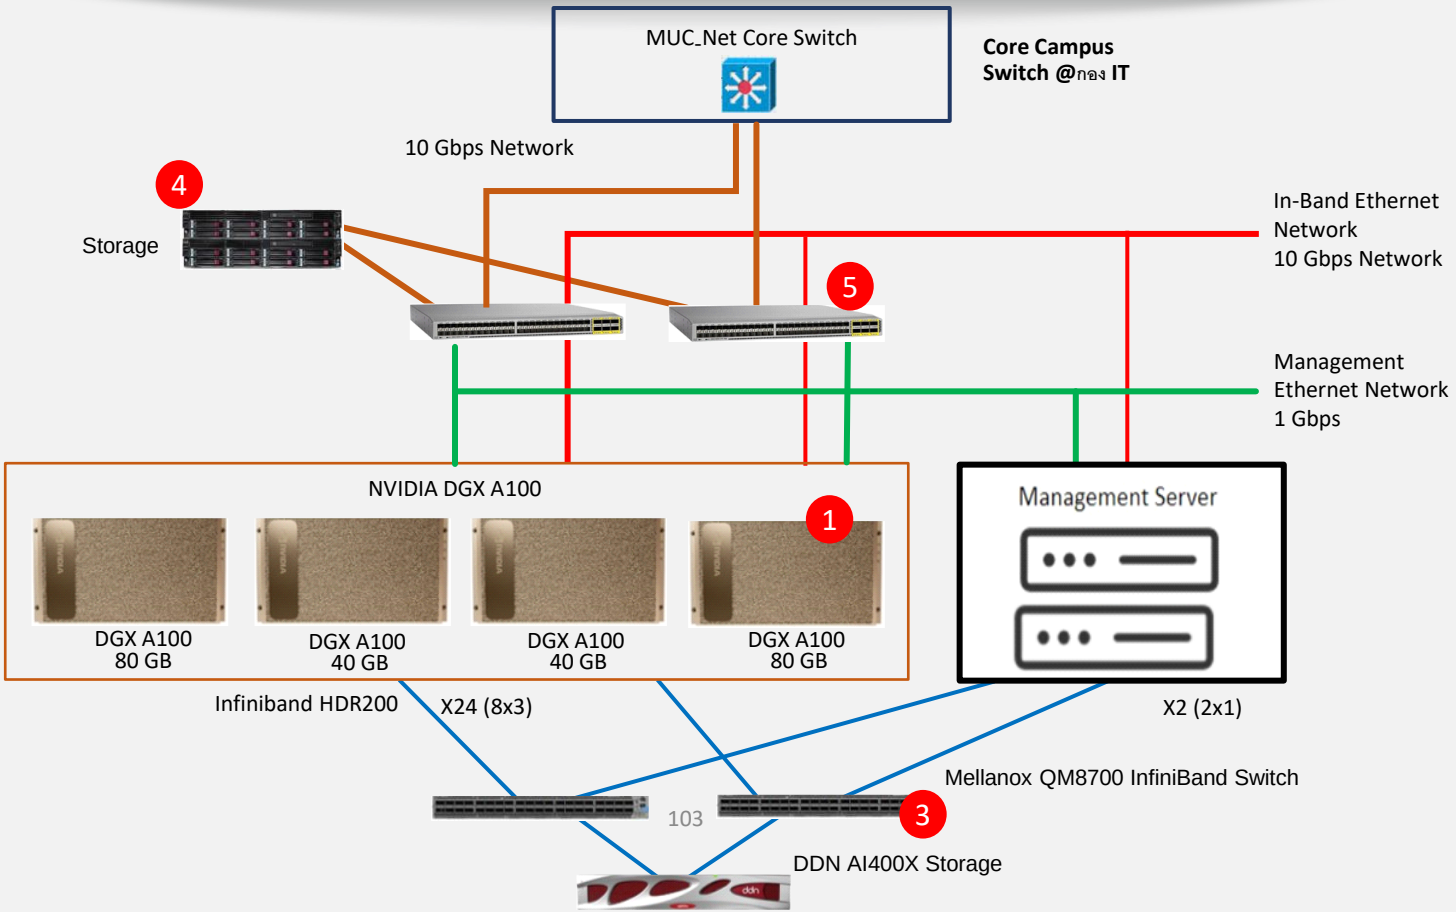

Thank you

References

1. Ethics of Artificial Intelligence | Internet Encyclopedia of Philosophy (utm.edu)
2. Ethics of Artificial Intelligence and Robotics (Stanford Encyclopedia of Philosophy)