Machine Learning and Robotics

Tushar B. Kute,
http://tusharkute.com

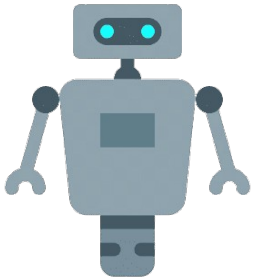

About me...

Tushar B. Kute

- Who: Researcher and Trainer
- Where: MITU Skillologies
- What: M.E. CSE
- When: 2005 (B.E. from COEP)
- Find me:
  - http://tusharkute.com
  - tushar@tusharkute.com
  - 7588594665

Artificial Intelligence

- Artificial intelligence (AI) refers to the simulation of human intelligence in machines that are programmed to think like humans and mimic their actions.
- The term may also be applied to any machine that exhibits traits associated with a human mind such as learning and problem-solving.

Artificial Intelligence

- Artificial intelligence is based on the principle that human intelligence can be defined in a way that a machine can easily mimic it and execute tasks, from the most simple to those that are even more complex.
- The goals of artificial intelligence include mimicking human cognitive activity.

Artificial Intelligence

Humans + Computers > Humans Alone

Humans + Computers > Computers Alone

Types

- “Strong” Artificial Intelligence
  - Computers thinking at a level that meets or surpasses people
  - Computers engaging in abstract reasoning & thinking
  - This is not what we have today
  - There is no evidence that we are close to Strong AI
- “Weak” Pattern-Based Artificial Intelligence
  - Computers solve problems by detecting useful patterns
  - Pattern-based AI is an Extremely powerful tool
  - Has been used to automate many processes today
  - Driving, language translation
  - This is the dominant mode of AI today

Major AI Approaches

- Two Major AI Techniques
  - Logic and Rules-Based Approach

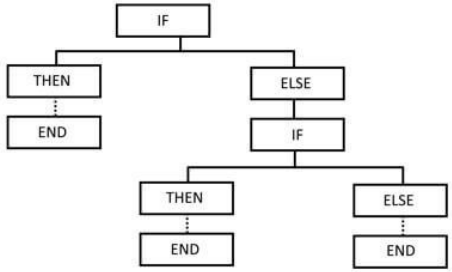

  - Machine Learning (Pattern-Based Approach)

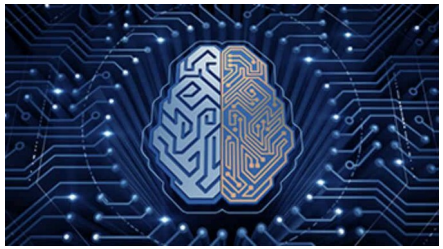

Rules based approach

- Logic and Rules-Based Approach
  - Representing processes or systems using logical rules
  - Top-down rules are created for computer
  - Computers reason about those rules
  - Can be used to automate processes
  - Example within law – Expert Systems
    - Turbotax
    - Personal income tax laws
    - Represented as logical computer rules
    - Software computes tax liability

Learning vs. Designing

- AI is a bigger concept to design intelligent machines that can simulate human thinking capability and behavior, whereas, machine learning is an application or subset of AI that allows machines to learn from data without being programmed explicitly.

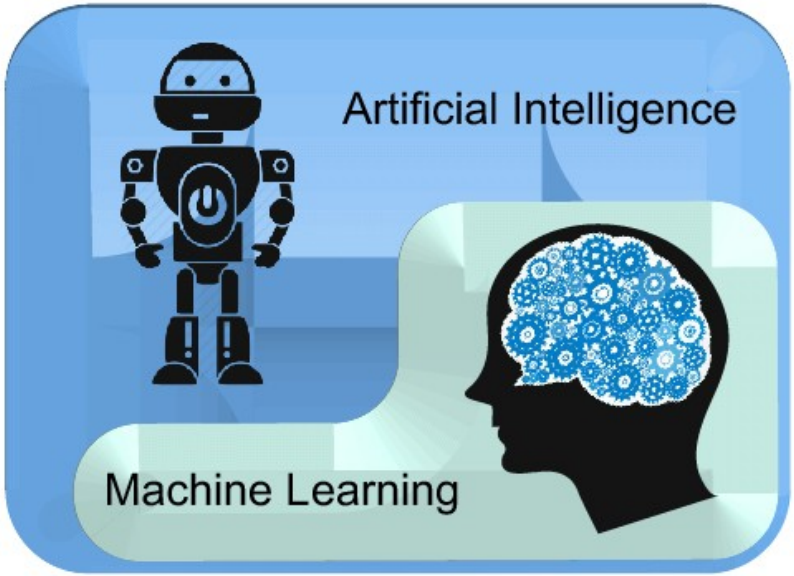

Technologies of next generation

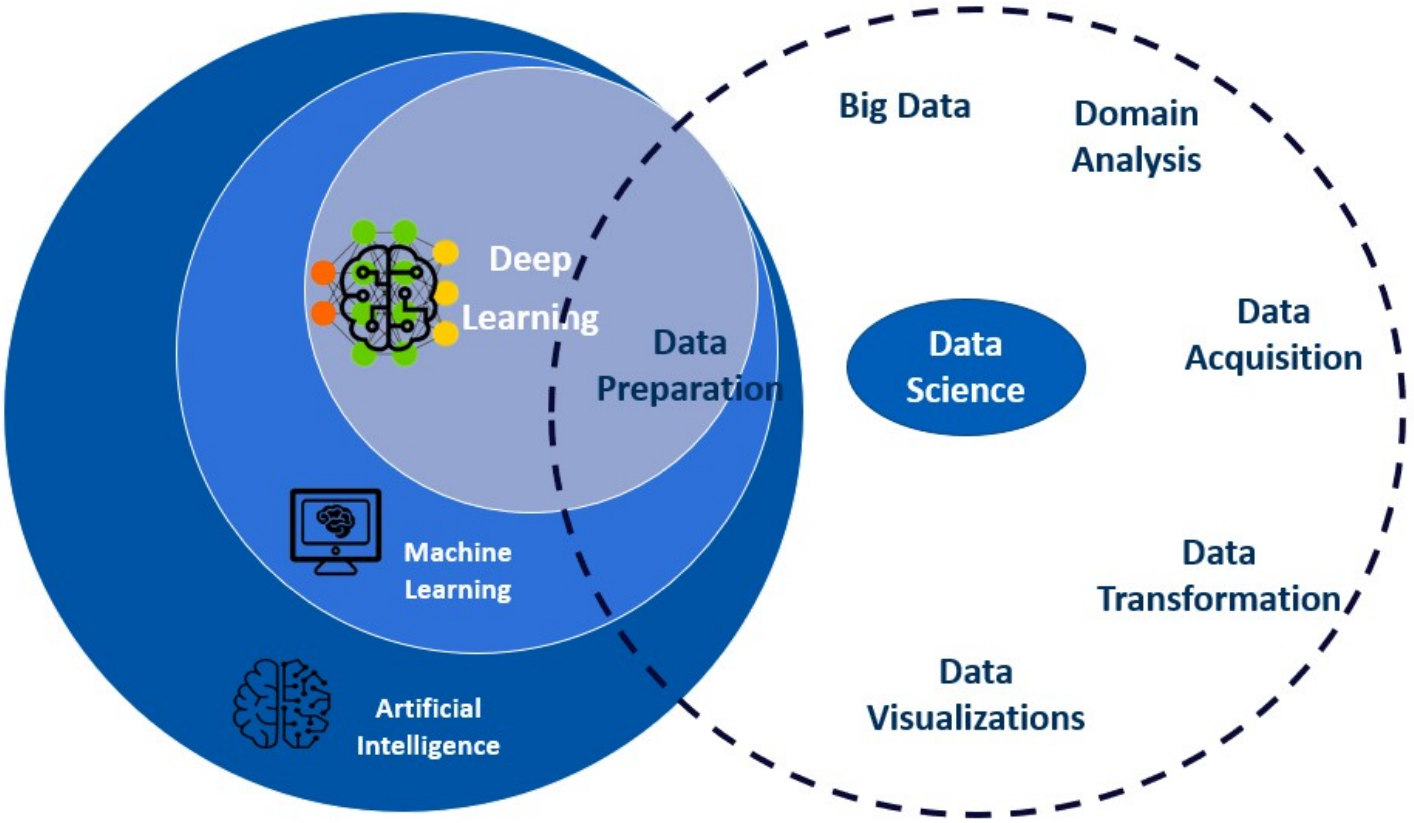

Technologies of next generation

Machine Learning vs AI

What is it???

It's a helicopter

I know how it flies

Data → Experience → Learning + Reasoning

Machine Learning

Artificial Intelligence

Machine Learning

- Machine learning is an application of artificial intelligence (AI) that provides systems the ability to automatically learn and improve from experience without being explicitly programmed. Machine learning focuses on the development of computer programs that can access data and use it learn for themselves.
- The process of learning begins with observations or data, such as examples, direct experience, or instruction, in order to look for patterns in data and make better decisions in the future based on the examples that we provide.
- The primary aim is to allow the computers learn automatically without human intervention or assistance and adjust actions accordingly.

Origins of Machine Learning

- The earliest databases recorded information from the observable environment.
- Astronomers recorded patterns of planets and stars; biologists noted results from experiments crossbreeding plants and animals; and cities recorded tax payments, disease outbreaks, and populations. Each of these required a human being to first observe and second, record the observation.
- Today, such observations are increasingly automated and recorded systematically in ever-growing computerized databases.

Machine Learning

Traditional Programming

Machine Learning

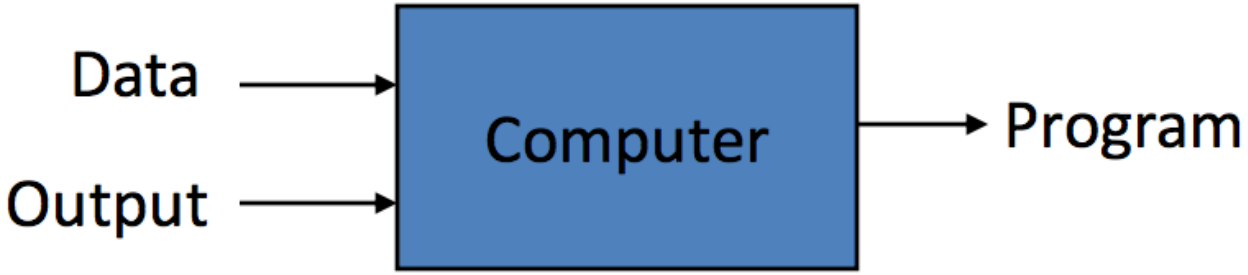

Machine Learning

- The field of study interested in the development of computer algorithms for transforming data into intelligent action is known as machine learning.

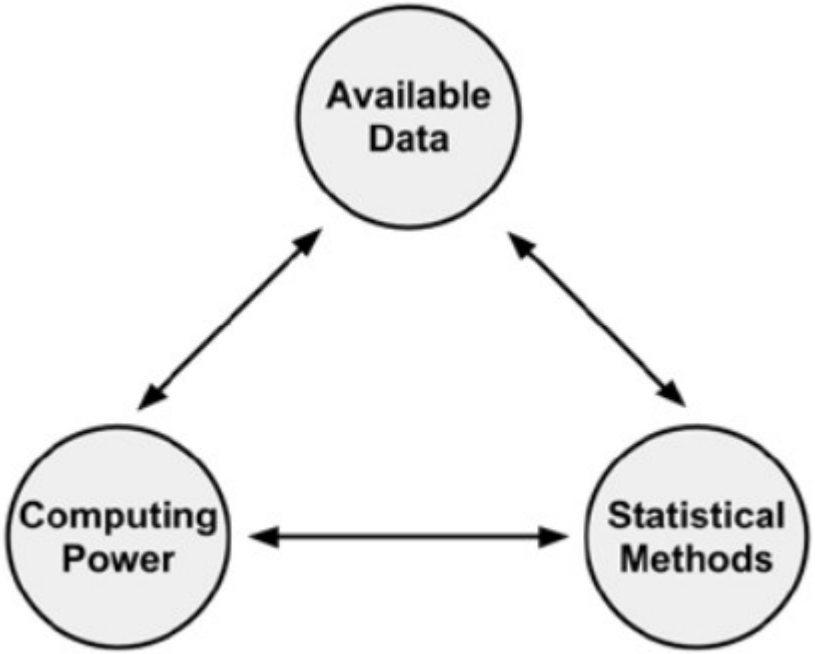

Uses and Abuses

- Predict the outcomes of elections
- Identify and filter spam messages from e-mail
- Foresee criminal activity
- Automate traffic signals according to road conditions
- Produce financial estimates of storms and natural disasters
- Examine customer churn
- Create auto-piloting planes and auto-driving cars
- Stock market predition
- Target advertising to specific types of consumers

Case Study

सकाळ

विद्यापीठात विद्यार्थ्यांचा 'एक्झिट पोल'

'रँडम फॉरेस्ट मॉडेल'नुसार युतीच राज्यात आघाडीवर

पुणे, ता. २१ : राज्यात भाजप आणि शिवसेना युती आघाडीवर असेल, असा अंदाज वर्तविणाऱ्या चाचण्यांचे कल (एक्झिट पोल) नुकतेच प्रसिद्ध झाले आहेत. सावित्रीबाई फुले पुणे विद्यापीठातील विद्यार्थ्यांनीही त्याला दुजोरा दिला आहे. भारतीय जनता पक्षाला १७ ते २३ आणि शिवसेनेला १६ ते २१ जागा मिळतील, असा अंदाज विद्यार्थ्यांनी 'रँडम फॉरेस्ट मॉडेल' पद्धत वापरून वर्तविला आहे. राष्ट्रवादी काँग्रेसला ३ ते ९ व काँग्रेसला १ ते ६ जागा मिळतील, असा अंदाज त्यांनी वर्तवला आहे.

विद्यापीठाच्या संख्याशास्त्र विभागातील एमएस्सी (द्वितीय वर्ष)

करणारे विनय तिवारी, आर. विश्वनाथ, शरद कोळसे या विद्यार्थ्यांनी सहायक प्राध्यापक डॉ. आकांक्षा काशीकर यांच्या मार्गदर्शनाखाली हा अंदाज दिला आहे.

निवडणूक आयोगाच्या संकेतस्थळावरून सर्वेक्षणासाठी लागणारी माहिती त्यांनी मिळविली. जनमानसाचा कल ओळखण्यासाठी 'सीएसडीएस-लोकनीती' सर्वेक्षण अहवालातून नोंदी घेतल्या. त्याचबरोबर सध्याच्या सरकारच्या कामगिरीबद्दल लोकांच्या प्रतिक्रिया, पंतप्रधानपदाच्या संभाव्य उमेदवारांची लोकप्रियता, मागील निवडणुकीतील आपले मत यंदा बदलू इच्छिणारे मतदार यांचा अभ्यास करण्यात आला. या अंदाजासाठी रँडम फॉरेस्ट मॉडेल वापरण्यापूर्वी २००९ आणि २०१४च्या निवडणुकांचे अंदाज पडताळून पाहण्यात आले. हे अंदाज प्रत्यक्ष निकालांशी पडताळून पाहिले असता, ते जवळपास ९६ टक्के जुळत असल्याचे निदर्शनास आले. म्हणूनच अभ्यासात माहितीच्या विश्लेषणासाठी या पद्धतीचा वापर करण्यात आला, असे डॉ. काशीकर यांनी सांगितले.

> संख्याशास्त्र आणि संगणकशास्त्र याची सांगड घालून आणि मशिन लर्निंगच्या साह्याने उपलब्ध माहितीचे विश्लेषण केले. संख्याशास्त्रातील अभ्यासाची वेगवेगळी मॉडेल वापरून १९७७ पासून ते आतापर्यंतच्या लोकसभा आणि विधानसभा निवडणुकीतील माहितीचा अभ्यास केला. त्यामुळे संख्याशास्त्राचा वापर करून वर्तविलेला अंदाज हा निवडणुकीच्या निकालांच्या जवळ जाणारा असेल.
> - शरद कोळसे, विद्यार्थी

Pune, Main
22/05/2019 Page No. 4

Recognizing patterns

- Pattern recognition is the automated recognition of patterns and regularities in data. It has applications in
  - statistical data analysis,
  - signal processing,
  - image analysis,
  - information retrieval,
  - bioinformatics,
  - data compression,
  - computer graphics and
  - machine learning.

How do machine learn ?

- A commonly cited formal definition of machine learning, proposed by computer scientist Tom M. Mitchell, says that a machine is said to learn if it is able to take experience and utilize it such that its performance improves up on similar experiences in the future.
- This definition is fairly exact, yet says little about how machine learning techniques actually learn to transform data into actionable knowledge.

Training a dataset

- The process of fitting a particular model to a dataset is known as training.
- Why is this not called learning? First, note that the learning process does not end with the step of data abstraction.
- Learning requires an additional step to generalize the knowledge to future data.
- Second, the term training more accurately describes the actual process undertaken when the model is fitted to the data.

Practical Machine Learning

X	Y	Z
5	2	14
8	5	22
4	8	14
9	2	20
7	1	15
7	8	23

Inputs

Output

Z = ?

---> ML Model

Practical Machine Learning

X	Y	Z	Pre	Error
5	2	14	12	-2
8	5	22	21	-1
4	8	14	16	+2
9	2	20	20	0
7	1	15	15	0
7	8	23	22	-1

Z = 2X + Y ---> ML Model

Prediction ---> X = 6 Y = 8 Z = ?

if 20 == 19: 95%

Training a dataset

Observations → **Data** → **Model**

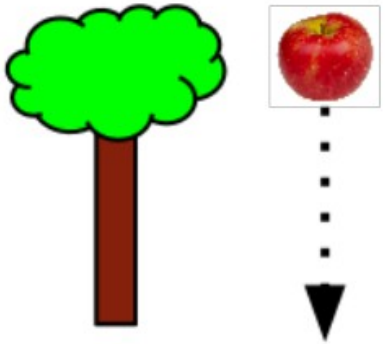

velocity	time
9.8	1
39.2	2
88.2	3
156.8	4
245	5

$$g = 9.8\ m/s^2$$

Types of Machine Learning

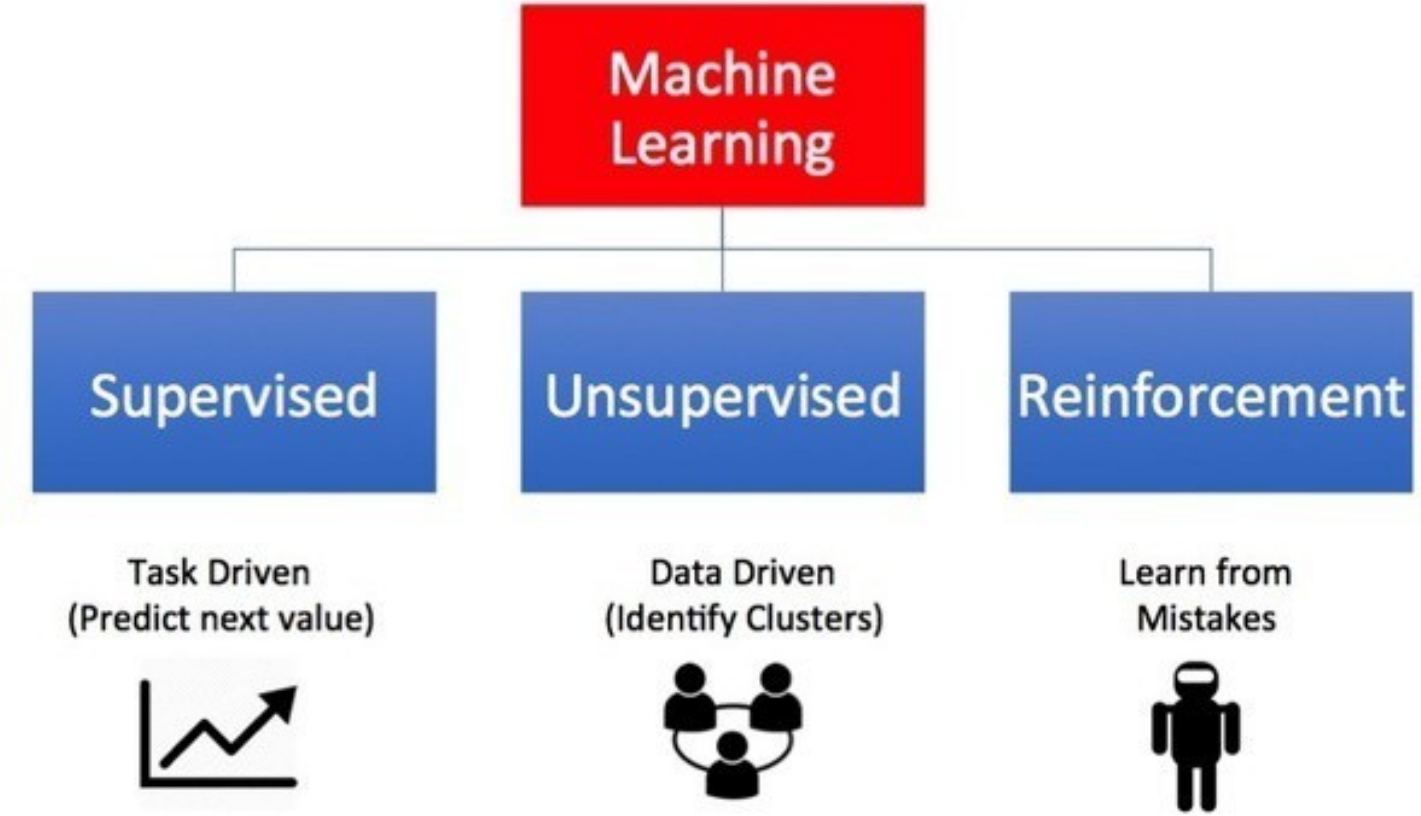

Supervised Machine Learning

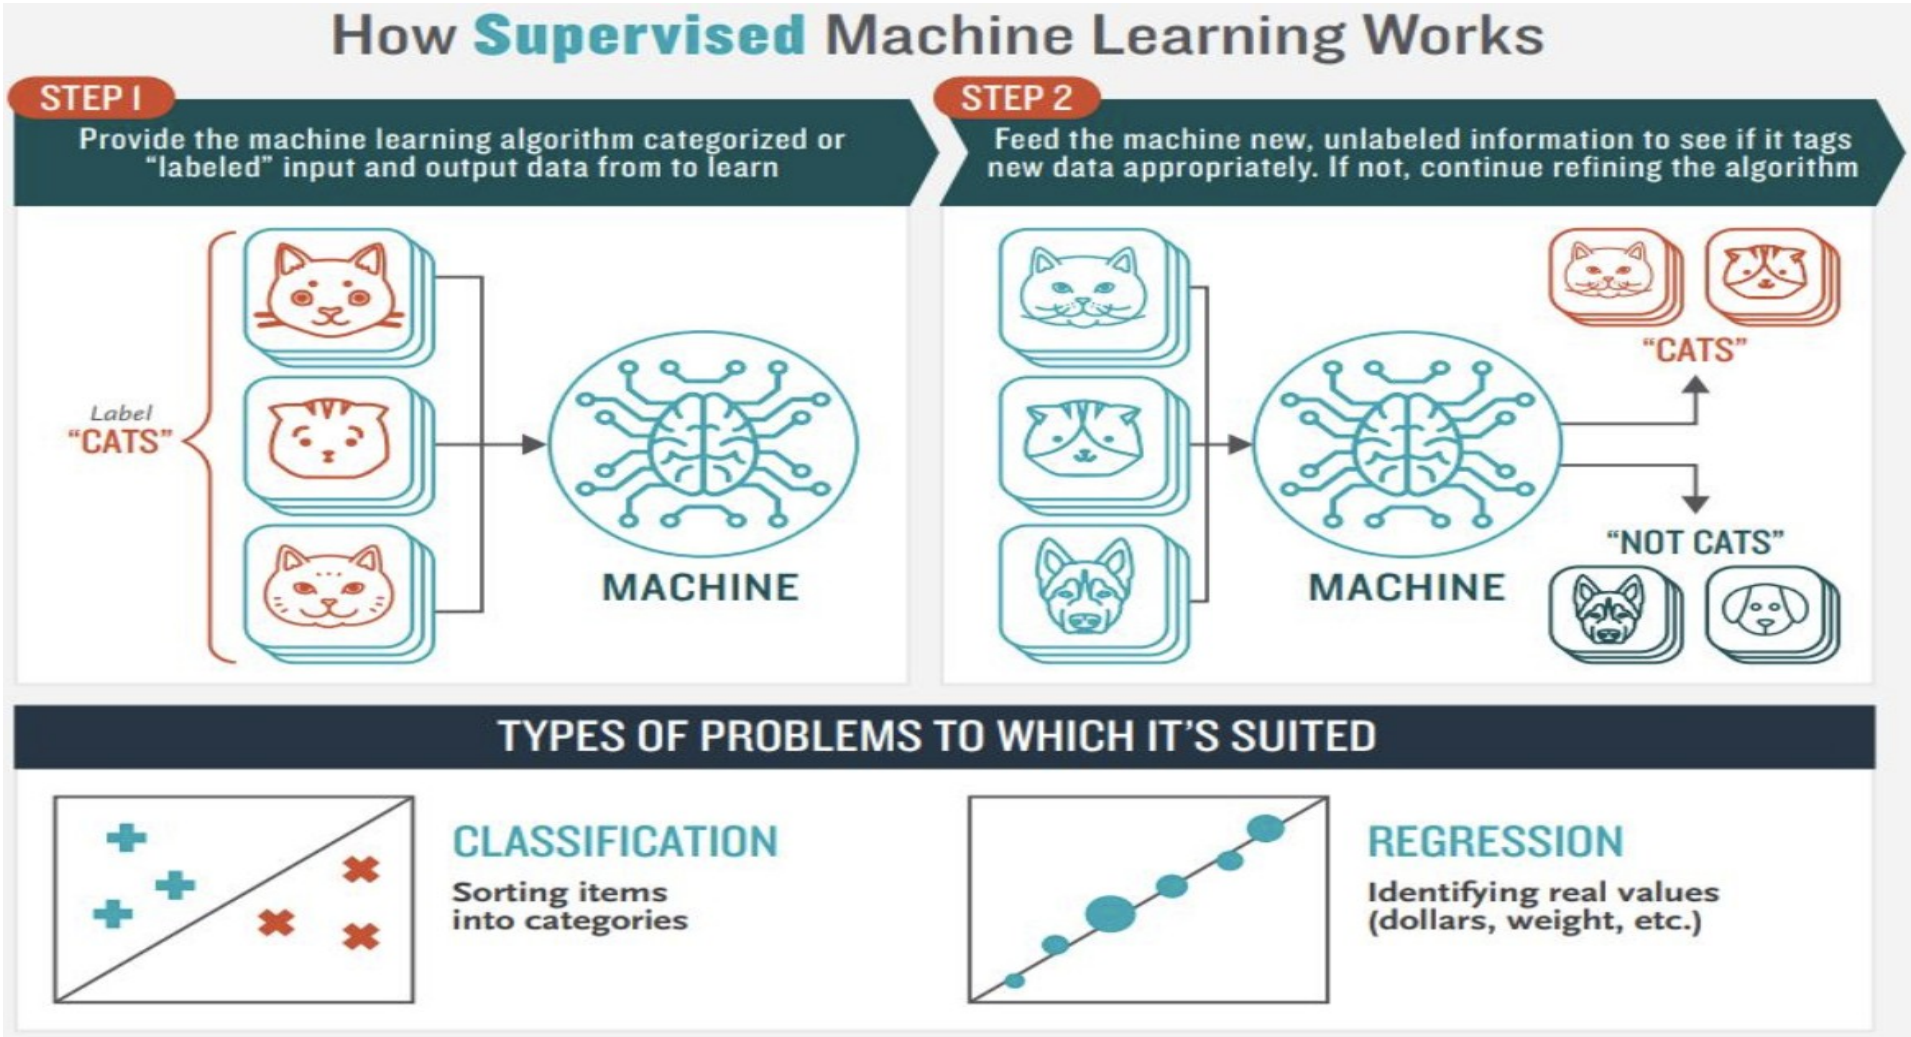

Unsupervised Machine Learning

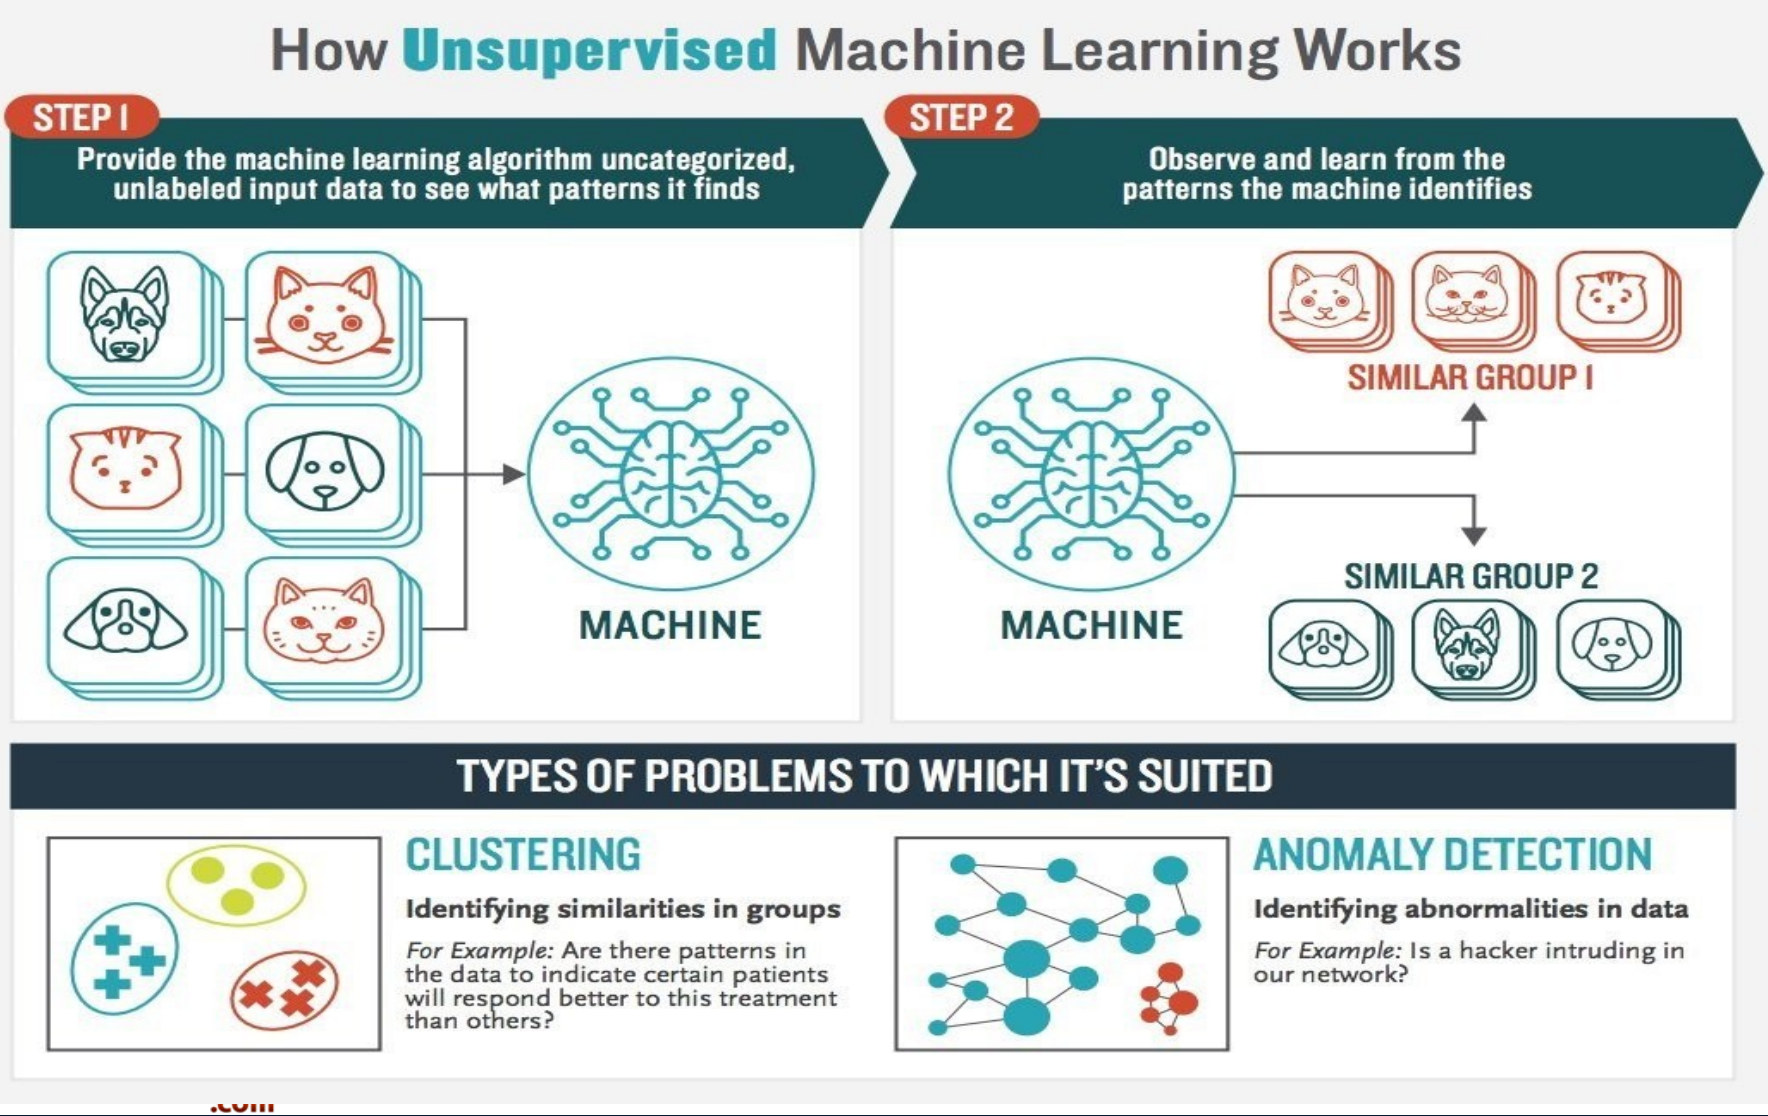

Typical Algorithms in Machine Learning

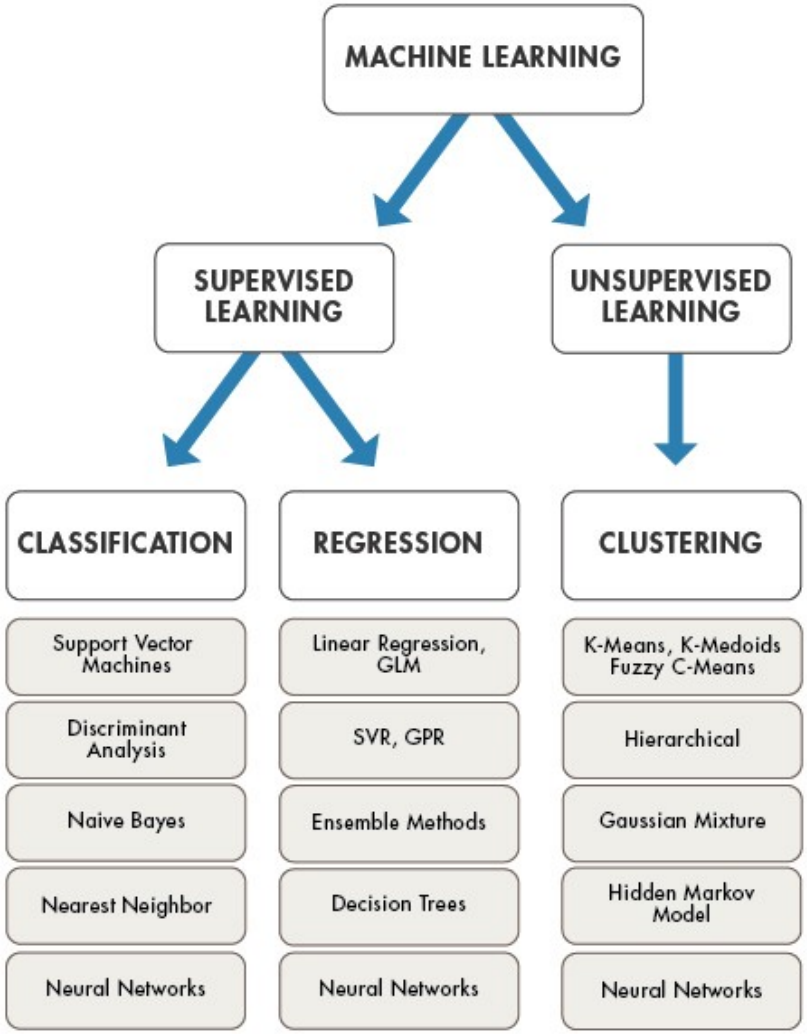

Intelligent Robots

- You can give it different names, an intelligent software module, an algorithm, a machine learning model, or a computer program that makes the intelligent decision for the integrated hardware interface function, without being explicitly programmed.”

Intelligent Robots: Examples

- 1- Tesla — is an excellent example of an intelligent robot that makes dynamic decisions based on cognitive power and leverages the learning and experience infused by a software program.
- 2- Drones — are autonomous and AI-based robots. They can mimic human intelligence in navigating their way to the target. They can also do all kinds of things without ever being programmed to do specific static actions.
- 3- Alexa — is another robot that performs specific actions based on human voice commands. Alexa provides an API (Application Programming Interface) based platform to integrate with any other pluggable interface.

Intelligent Robots

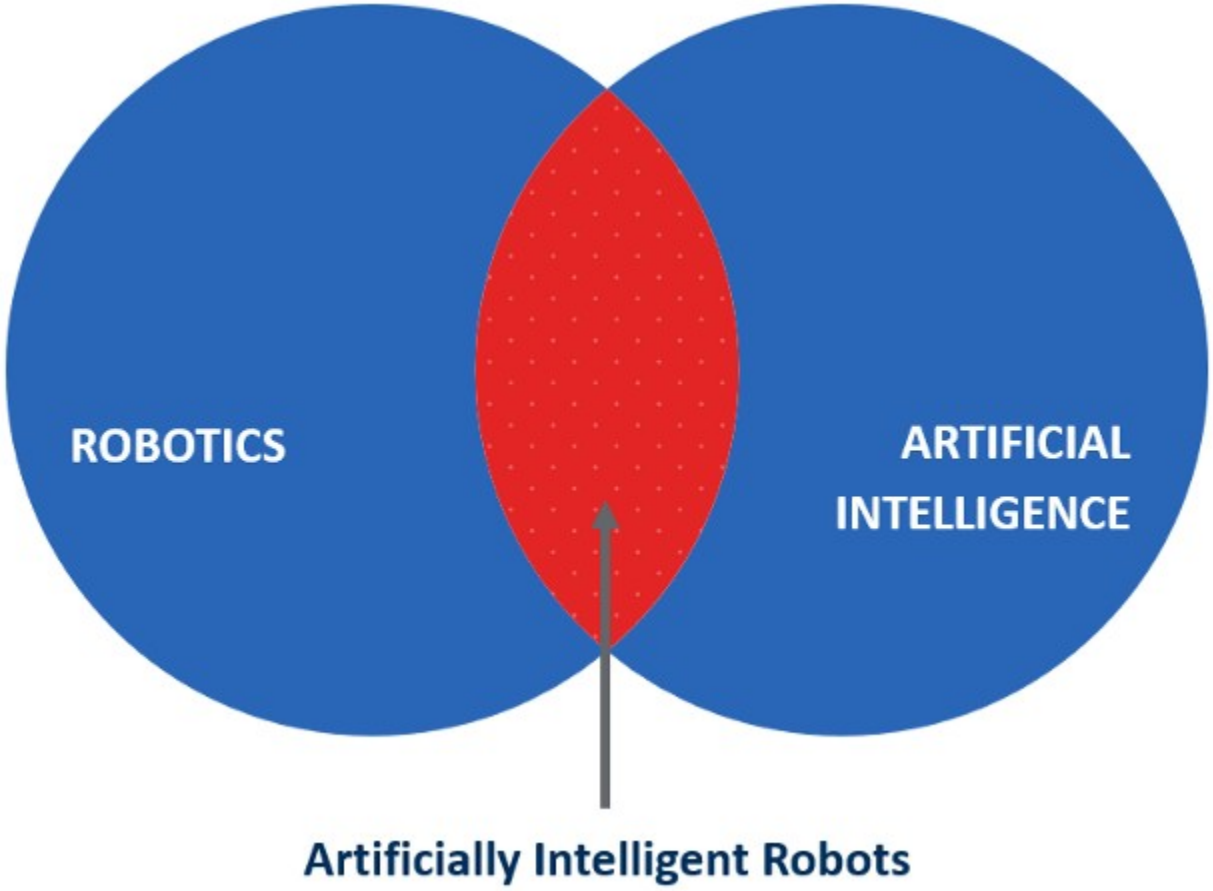

Scope of ML in Robotics

- Vision – AI is helping robots detect items they've never seen before and recognize objects with far greater detail.
- Grasping – robots are also grasping items they've never seen before with AI and machine learning helping them determine the best position and orientation to grasp an object.
- Motion Control – machine learning helps robots with dynamic interaction and obstacle avoidance to maintain productivity.
- Data – AI and machine learning both help robots understand physical and logistical data patterns to be proactive and act accordingly.

ML & Robotics: Computer Vision

- Though related, some would argue the correct term is machine vision or robot vision rather than computer vision, because “robots seeing” involves more than just computer algorithms; engineers and roboticists also have to account for camera hardware that allow robots to process physical data.
- Robot vision is very closely linked to machine vision, which can be given credit for the emergence of robot guidance and automatic inspection systems.
- The slight difference between the two may be in kinematics as applied to robot vision, which encompasses reference frame calibration and a robot’s ability to physically affect its environment.

ML & Robotics: Imitation Learning

- Imitation learning is closely related to observational learning, a behavior exhibited by infants and toddlers.
- Imitation learning is also an umbrella category for reinforcement learning, or the challenge of getting an agent to act in the world so as to maximize its rewards.
- Bayesian or probabilistic models are a common feature of this machine learning approach.
- The question of whether imitation learning could be used for humanoid-like robots was postulated as far back as 1999.

ML & Robotics: Self Supervised Learning

- Self-supervised learning approaches enable robots to generate their own training examples in order to improve performance; this includes using a priori training and data captured close range to interpret “long-range ambiguous sensor data.”
- It’s been incorporated into robots and optical devices that can detect and reject objects (dust and snow, for example); identify vegetables and obstacles in rough terrain; and in 3D-scene analysis and modeling vehicle dynamics

ML & Robotics: Self Supervised Learning

- Watch-Bot is a concrete example, created by researchers from Cornell and Stanford, that uses a 3D sensor (a Kinect), a camera, laptop and laser pointer to detect ‘normal human activity’, which are patterns that it learns through probabilistic methods.
- Watch-Bot uses a laser pointer to target the object as a reminder (for example, the milk that was left out of the fridge).
- In initial tests, the bot was able to successfully remind humans 60 percent of time (it has no conception of what it’s doing or why), and the researchers expanded trials by allowing its robot to learn from online videos (called project RoboWatch).

ML & Robotics: Self Supervised Learning

- Other examples of self-supervised learning methods applied in robotics include a road detection algorithm in a front-view monocular camera with a road probabilistic distribution model (RPDM) and fuzzy support vector machines (FSVMs), designed at MIT for autonomous vehicles and other mobile on-road robots.
- Autonomous learning, which is a variant of self-supervised learning involving deep learning and unsupervised methods, has also been applied to robot and control tasks.

ML & Robotics: Medical Technologies

- An assistive robot (according to Stanford's David L. Jaffe) is a device that can sense, process sensory information, and perform actions that benefit people with disabilities and seniors (though smart assistive technologies also exist for the general population, such as driver assistance tools). Movement therapy robots provide a diagnostic or therapeutic benefit.
- Both of these are technologies that are largely (and unfortunately) still confined to the lab, as they're still cost-prohibitive for most hospitals in the U.S. and abroad.

ML & Robotics: Medical Technologies

- Early examples of assistive technologies included the DeVAR, or desktop vocational assistant robot, developed in the early 1990s by Stanford and the Palo Alto Veterans Affairs Rehabilitation Research and Development.
- More recent examples of machine learning-based robotic assistive technologies are being developed that include combining assistive machines with more autonomy, such as the MICO robotic arm (developed at Northwester University) that observes the world through a Kinect Sensor.

ML & Robotics: Multi Agent Learning

- Coordination and negotiation are key components of multi-agent learning, which involves machine learning-based robots (or agents – this technique has been widely applied to games) that are able to adapt to a shifting landscape of other robots/agents and find “equilibrium strategies.”
- Examples of multi-agent learning approaches include no-regret learning tools, which involve weighted algorithms that “boost” learning outcomes in multi-agent planning, and learning in market-based, distributed control systems.

ML & Robotics: Multi Agent Learning

- A more concrete example is an algorithm for distributed agents or robots created by researchers from MIT's Lab for Information and Decision Systems in late 2014.
- Robots collaborated to build a better and more inclusive learning model than could be done with one robot (smaller chunks of information processed and then combined), based on the concept of exploring a building and its room layouts and autonomously building a knowledge base.

ML & Robotics: Kinematics

- Kinematics – Branch of classical mechanics which describes the motion of points (alternatively “particles”), bodies (objects), and systems of bodies without consideration of the masses of those objects nor the forces that may have caused the motion; often referred to as “geometry of motion”.

ML & Robotics: Bayesian Models

- Bayesian models – Method of statistical inference that casts statistical problems in the framework of decision making.
- It entails formulating subjective prior probabilities to express pre-existing information, careful modeling of the data structure, checking and allowing for uncertainty in model assumptions, formulating a set of possible decisions and a utility function to express how the value of each alternative decision is affected by the unknown model parameters.

ML & Robotics: Inverse Optimal Control

- Inverse optimal control – Also known as inverse reinforcement learning, it’s the problem of recovering an unknown reward function in a Markov decision process from expert demonstrations of the optimal policy.

ML & Robotics: SVM

- Support vector machines – Also called support vector networks, SVMs are supervised learning models with associated learning algorithms that analyze data used for classification and regression analysis.

Reference: https://emerj.com

Where to focus ?

- Beginning a career in Machine Learning obviously involves a lot of Self-learning, which can be broadly put in two categories:
  - Theory (Math, Computer Science, Statistics, Operations Research, other ML theory)
  - Practice (Programming in Python, R ; Working with databases and big data technologies; Practical aspects of Data wrangling and visualization, Model building and validation, etc. which come from doing and not reading)

Core Skill-set expected

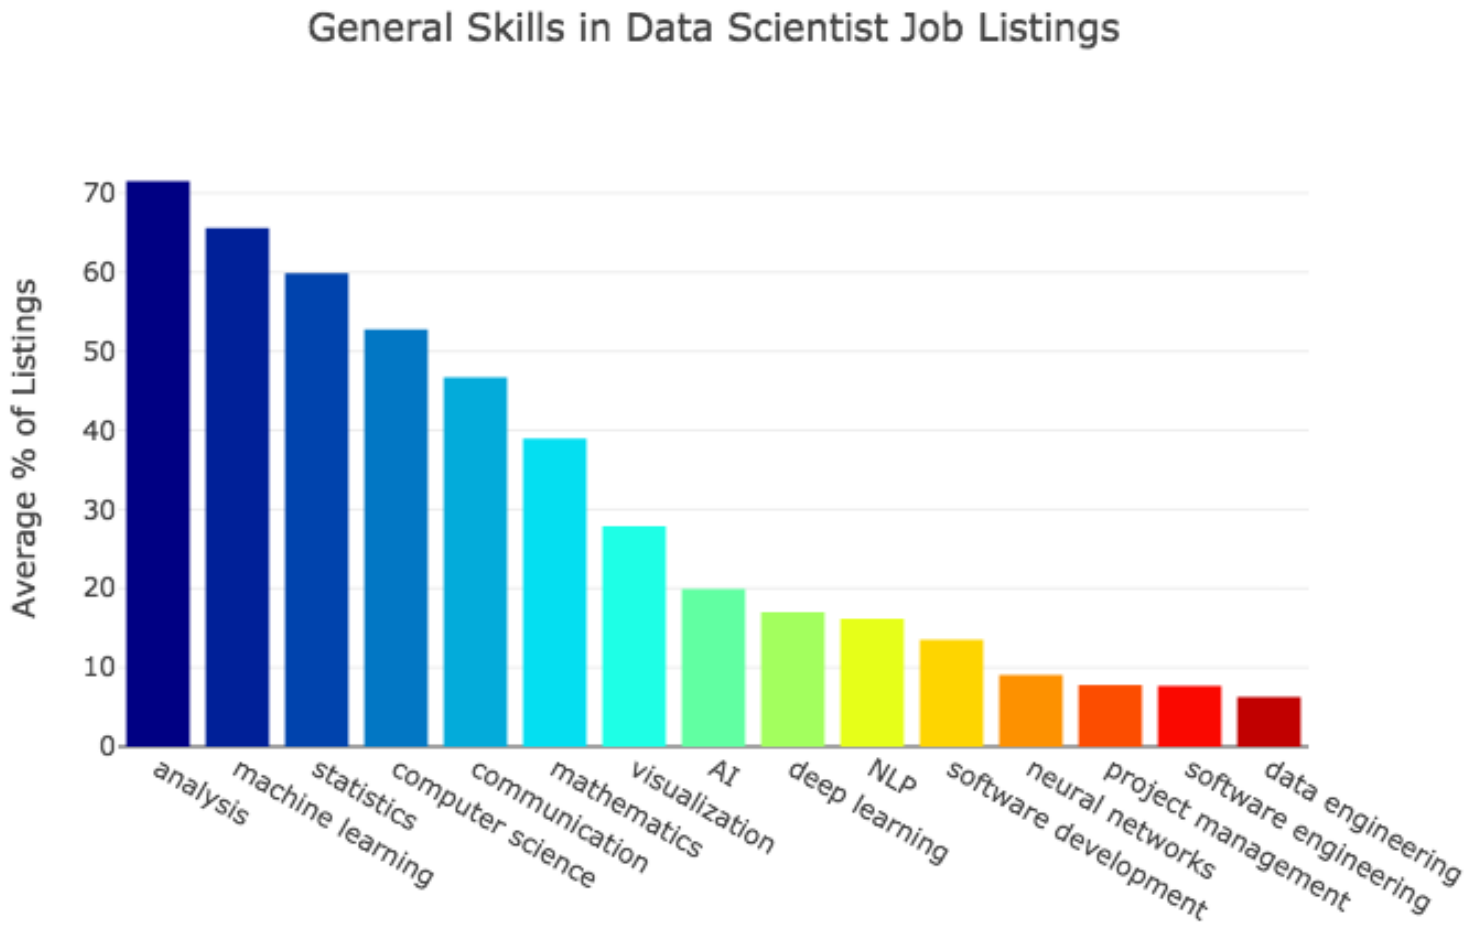

Reference: kdnuggets.com

Differentiating with skill-sets

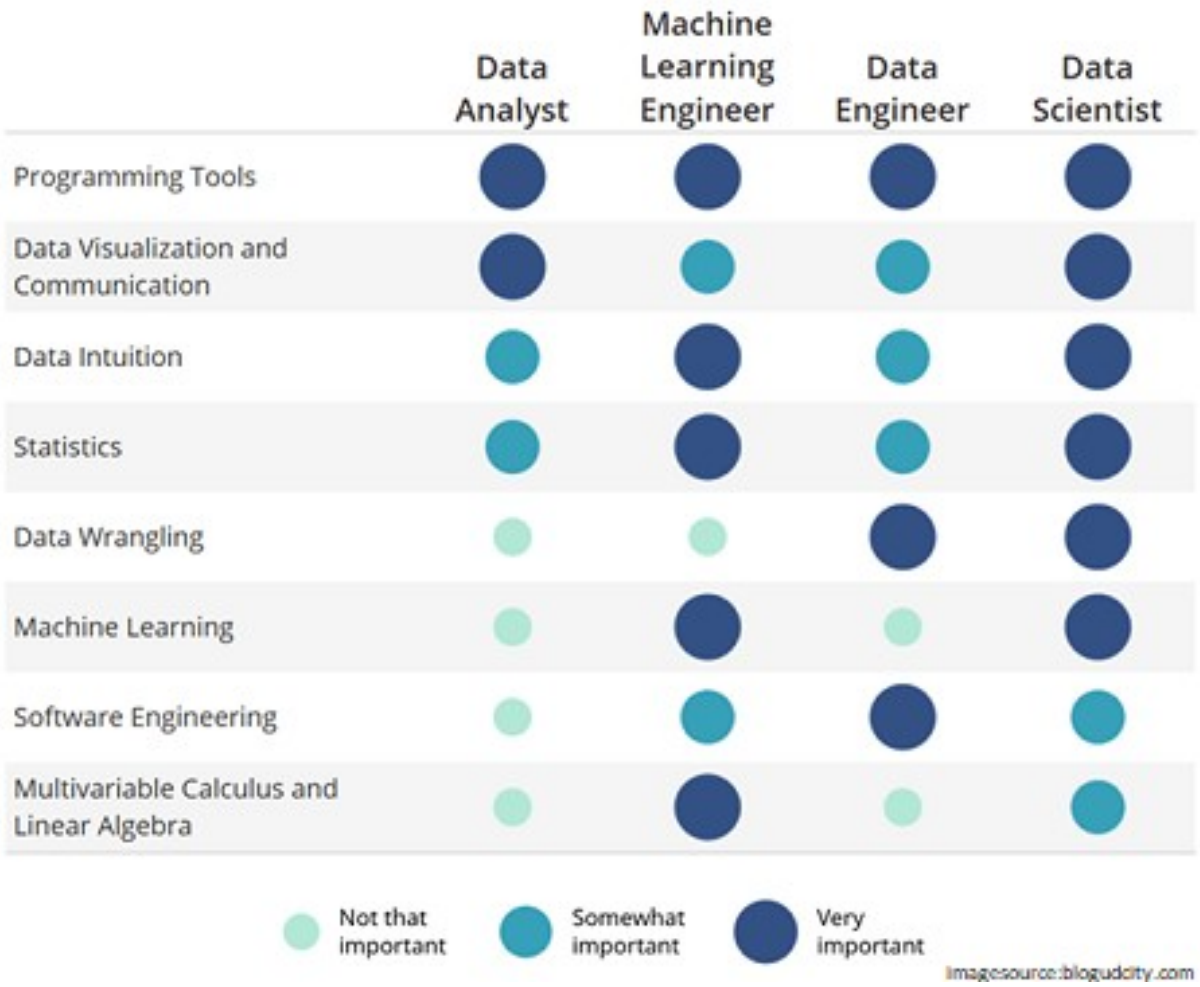

Average Annual Salary

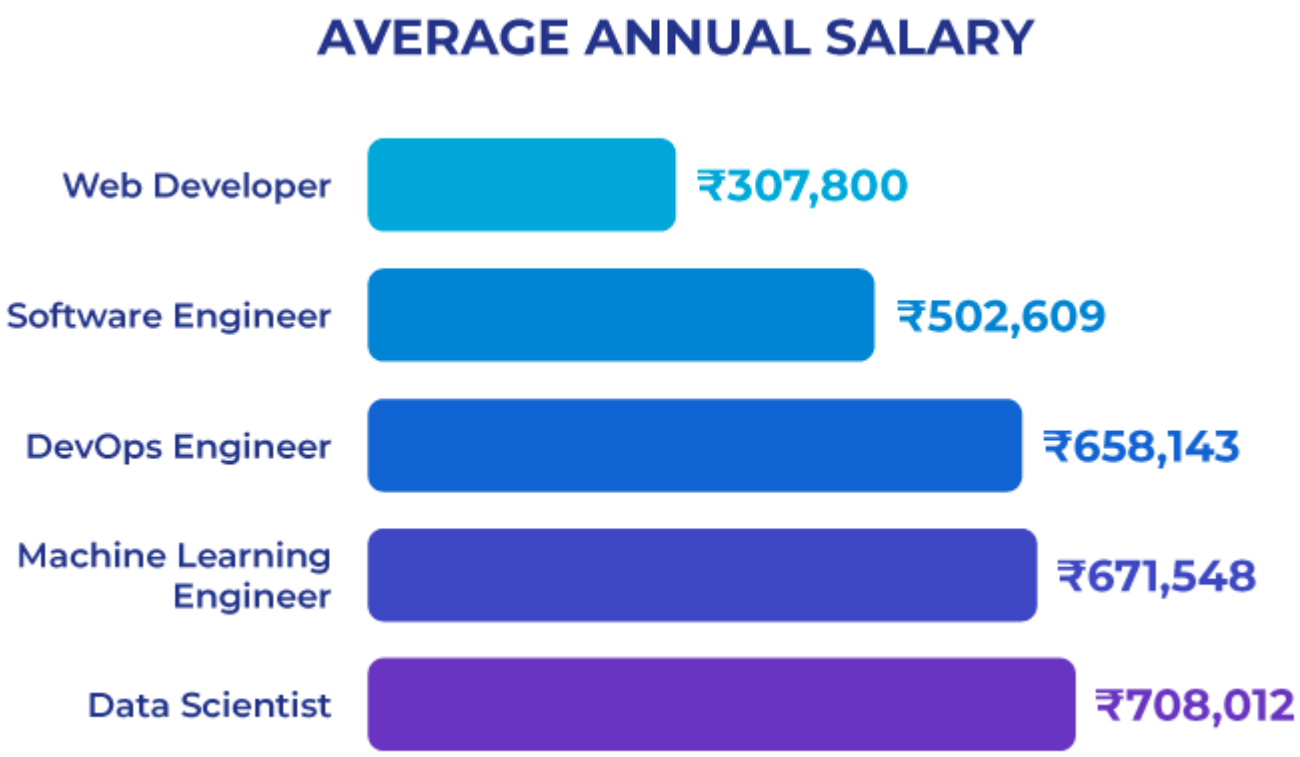

Conclusion

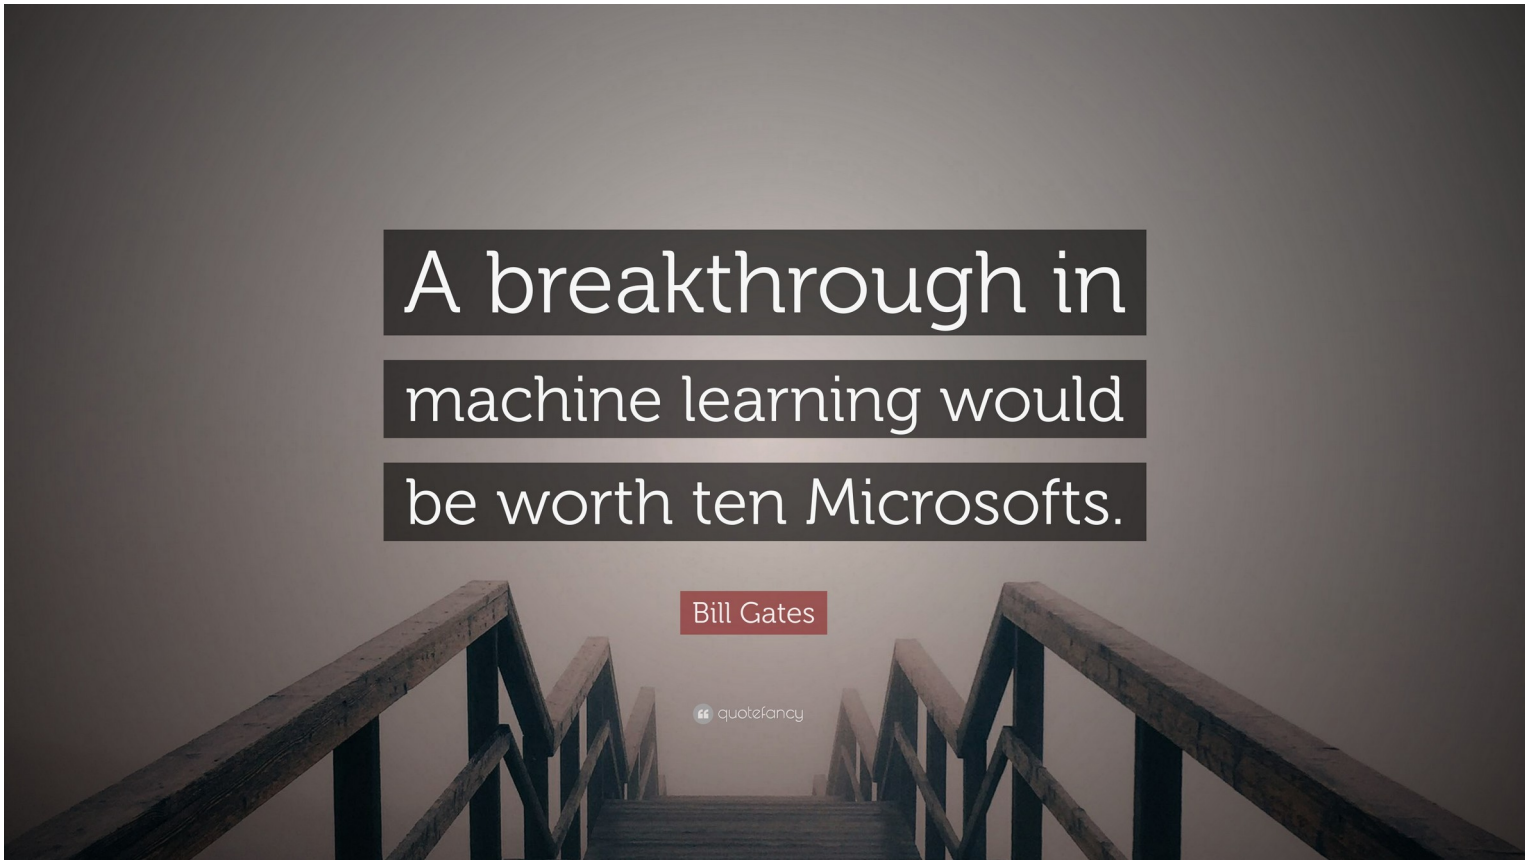
A breakthrough in
machine learning would
be worth ten Microsofts.
Bill Gates
quotefancy

Useful web resources

- www.mitu.co.in
- www.pythonprogramminglanguage.com
- www.scikit-learn.org
- www.towardsdatascience.com
- www.medium.com
- www.analyticsvidhya.com
- www.kaggle.com
- www.stephacking.com
- www.github.com

Thank you

This presentation is created using LibreOffice Impress 6.4.7.2, can be used freely as per GNU General Public License

@mitu_skillologies

/mITuSkillologies

@mitu_group

/company/mitu-skillologies

MITUSkillologies

Web Resources
https://mitu.co.in
http://tusharkute.com

contact@mitu.co.in

tushar@tusharkute.com