



INTRODUCTION TO AI – AI FOR EVERYONE

Legal Notices and Disclaimers

This presentation is for informational purposes only. INTEL MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at [intel.com](https://www.intel.com).

This sample source code is released under the [Intel Sample Source Code License Agreement](#).

Intel, the Intel logo, neon, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2018, Intel Corporation. All rights reserved.

Outline: What students are expected to Learn

1. What is AI
2. Background of AI on PC
3. Machine Learning
4. Deep Learning
5. Summary

THE DELUGE OF DATA

DAILY BY 2020


AVERAGE INTERNET USER **1.5 GB**

AUTONOMOUS VEHICLE **4 TB**

CONNECTED AIRPLANE **5 TB**

SMART FACTORY **1 PB**

CLOUD VIDEO PROVIDER **750 PB**



**BUSINESS
INSIGHTS**

**OPERATIONAL
INSIGHTS**

**SECURITY
INSIGHTS**

Source: Amalgamation of analyst data and Intel analysis.

AI WILL TRANSFORM



CONSUMER

Smart Assistants
Chatbots
Search
Personalization
Augmented Reality
Robots



HEALTH

Enhanced Diagnostics
Drug Discovery
Patient Care
Research
Sensory Aids



FINANCE

Algorithmic Trading
Fraud Detection
Research
Personal Finance
Risk Mitigation



RETAIL

Support Experience
Marketing
Merchandising
Loyalty
Supply Chain
Security



GOVERNMENT

Defense
Data Insights
Safety & Security
Resident Engagement
Smarter Cities



ENERGY

Oil & Gas Exploration
Smart Grid
Operational Improvement
Conservation



TRANSPORT

In-Vehicle Experience
Automated Driving
Aerospace
Shipping
Search & Rescue



INDUSTRIAL

Factory Automation
Predictive Maintenance
Precision Agriculture
Field Automation

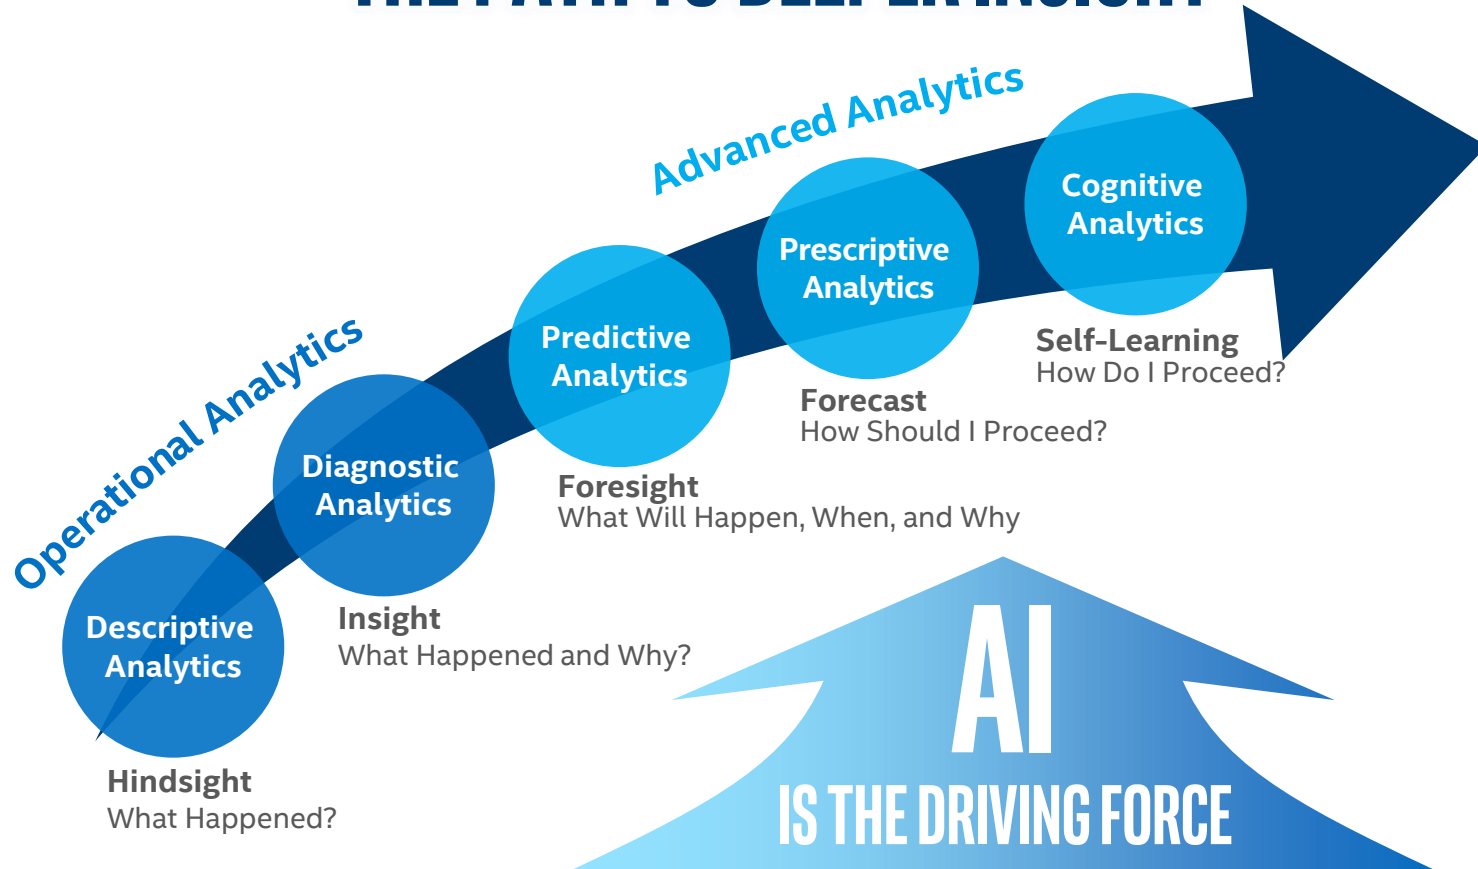


OTHER

Advertising
Education
Gaming
Professional & IT Services
Telco/Media
Sports

Source: Intel forecast

THE PATH TO DEEPER INSIGHT



AI ADOPTION IS NASCENT

According to a recent Gartner survey...

46%

of Chief Information Officers (CIOs) have developed plans to implement AI, but **only**

4%

have implemented AI so far.

THE AI LIFECYCLE

1. Define the Challenge

2. Approach

Team breaks down the defined business problem into workable steps to translate the right data to achieve results

3. Expertise

A team of management sponsors, data scientists, data engineers, solution architects, and domain experts identifies the right data and works to translate the data to achieve results

4. Philosophy

Team embraces fail-fast continuous improvement practices to evaluate their success in translating data to achieve results

7. Organization

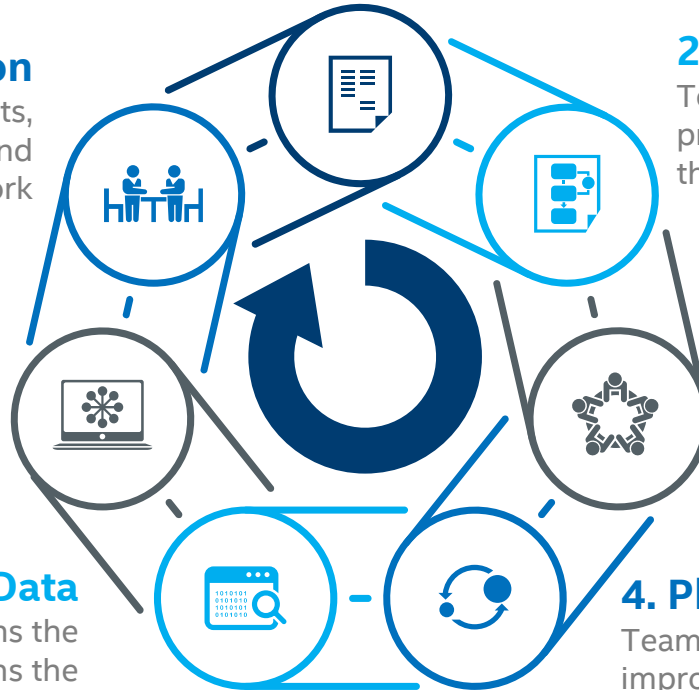
Organization embraces data insights, sponsors properly resourced teams, and prioritizes analytic development work


6. Infrastructure

Organization secures hardware and software infrastructure that supports data processing in a timely manner

5. Source Data

Team understands and obtains the right data that explains the business problem to achieve results





WHICH APPROACH IS RIGHT?

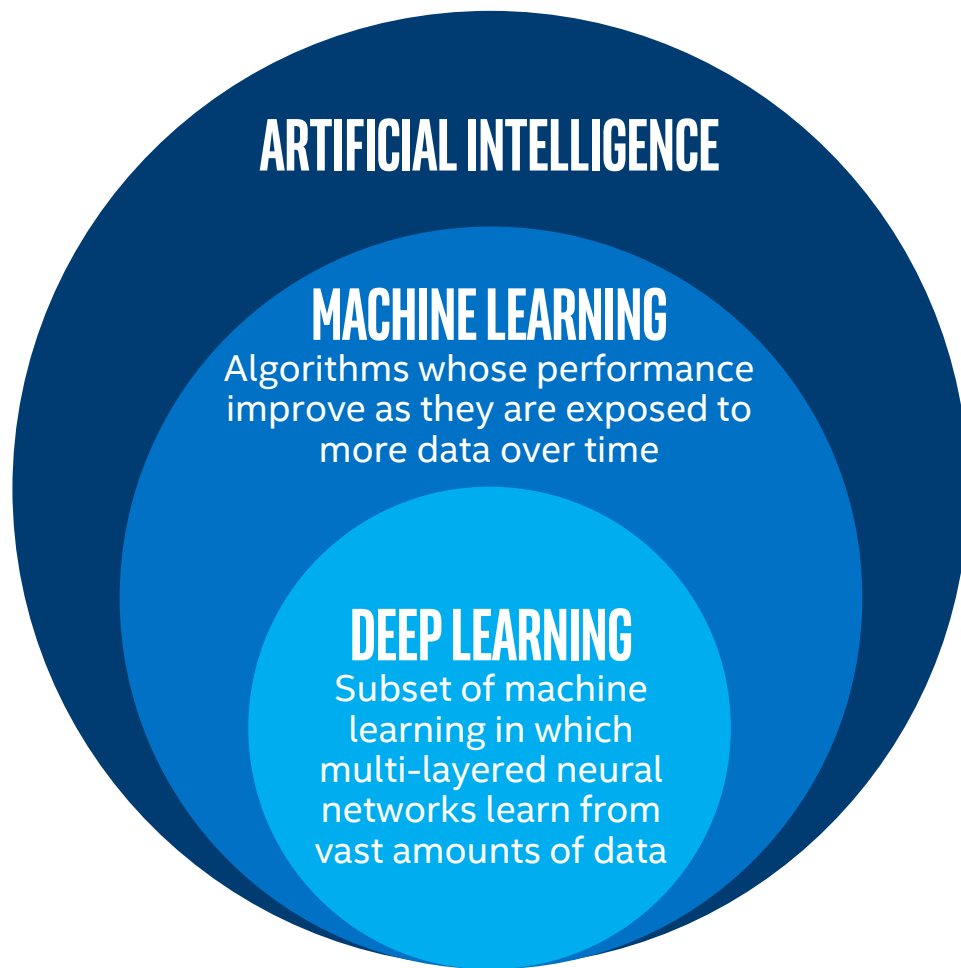
A large **manufacturer** uses data to improve their operations, with each challenge using a different approach to deliver maximum business value at the lowest possible cost

CHALLENGE	BEST APPROACH	APPROACH	ANSWER
How many widgets should we manufacture?	Analyze historical supply/demand	Analytics/ Business Intelligence	10,000
What will our yield be?	Algorithm that correlates many variables to yield	Statistical/ Machine Learning	At current conditions, yield will be at 90% with 10% loss expected
Which widgets have visual defects?	Algorithm that learns to identify defects in images	Deep Learning	Widget 1003, Widget 1094 ...

LEARN
MORE IN
THE NEXT
SLIDES

ARTIFICIAL INTELLIGENCE

is the ability of machines to learn from experience, without explicit programming, in order to perform cognitive functions associated with the human mind



Artificial Intelligence

- “A branch of computer science dealing with the simulation of intelligent behavior in computers.” (Merriam-Webster)
- “A program that can sense, reason, act, and adapt.” (Intel)
- “Colloquially, the term ‘artificial intelligence’ is applied when a machine mimics ‘cognitive’ functions that humans associate with other human minds, such as ‘learning’ and ‘problem solving’.” (Wikipedia)

Background of AI on PC


- AI has increased significantly in the last 5 years with the availability of large data sources, growth in compute engines and modern algorithms development. AI on PC is propelling new technologies into all parts of modern life. Central to this story is that the PCs are now well capable of applying AI technologies to varied usages from computer vision to identification, classification to natural language processing.
- As a data scientist Andrew Ng noted, AI is the next electricity:

"Just as electricity transformed almost everything 100 years ago, today I actually have a hard time thinking of an industry that I don't think AI will transform in the next several years."

- While AI usage in the cloud continues to grow, there is an inclination to perform AI inference on the PCs driven by the need for lower latency, persistent availability, lower costs and addressing privacy concerns. We are moving to the day that devices from phones to PCs and embedded edge devices all will have AI embedded in them.

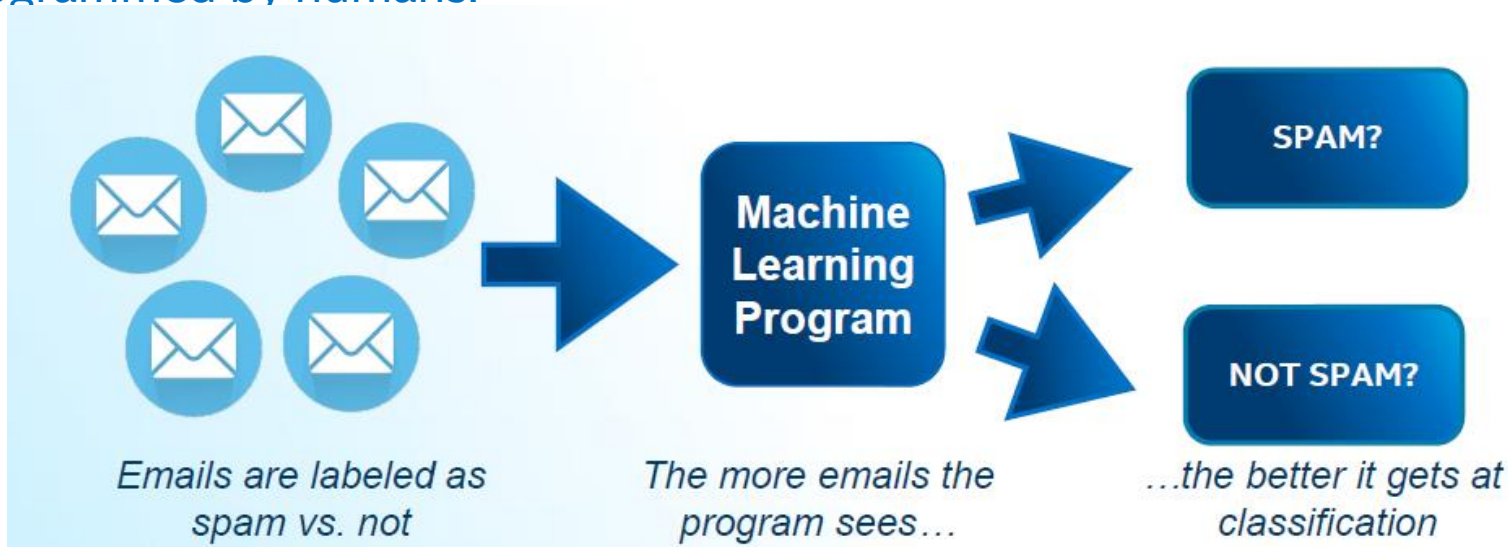
Machine Learning

“The study and construction of programs that are *not explicitly programmed*, but learn patterns as they are exposed to more data over time.” (Intel)



Machine Learning

These programs learn from repeatedly seeing data, rather than being explicitly programmed by humans.



Machine Learning Terminology

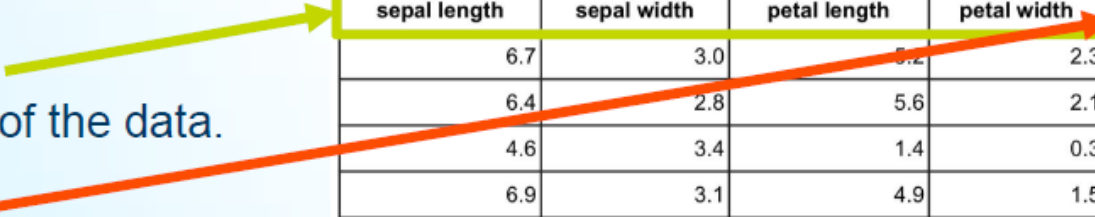
This example is learning to classify a species from a set of measurement features.

Features:

Attributes of the data.

Target:

Column to be predicted.

A diagram with two arrows originates from the text on the left. A yellow arrow points from 'Features:' to the first four columns of the table (sepal length, sepal width, petal length, petal width). An orange arrow points from 'Target:' to the 'species' column.

sepal length	sepal width	petal length	petal width	species
6.7	3.0	5.2	2.3	virginica
6.4	2.8	5.6	2.1	virginica
4.6	3.4	1.4	0.3	setosa
6.9	3.1	4.9	1.5	versicolor
4.4	2.9	1.4	0.2	setosa
4.8	3.0	1.4	0.1	setosa
5.9	3.0	5.1	1.8	virginica
5.4	3.9	1.3	0.4	setosa
4.9	3.0	1.4	0.2	setosa
5.4	3.4	1.7	0.2	setosa

Two Main Types of Machine Learning

	Dataset	Goal	Example
Supervised Learning	Has a target column	Make predictions	Fraud detection
Unsupervised Learning	Does not have a target column	Find structure in the data	Customer segmentation

Machine Learning Example – fraud detection

- Suppose you wanted to identify fraudulent credit card transactions.
- You could define features to be:
 - Transaction time
 - Transaction amount
 - Transaction location
 - Category of purchase
- The algorithm could learn what feature combinations suggest unusual activity.



Machine Learning Limitations

- Suppose you wanted to determine if an image is a dog.
- What features would you use?
- This is where Deep Learning can come in.

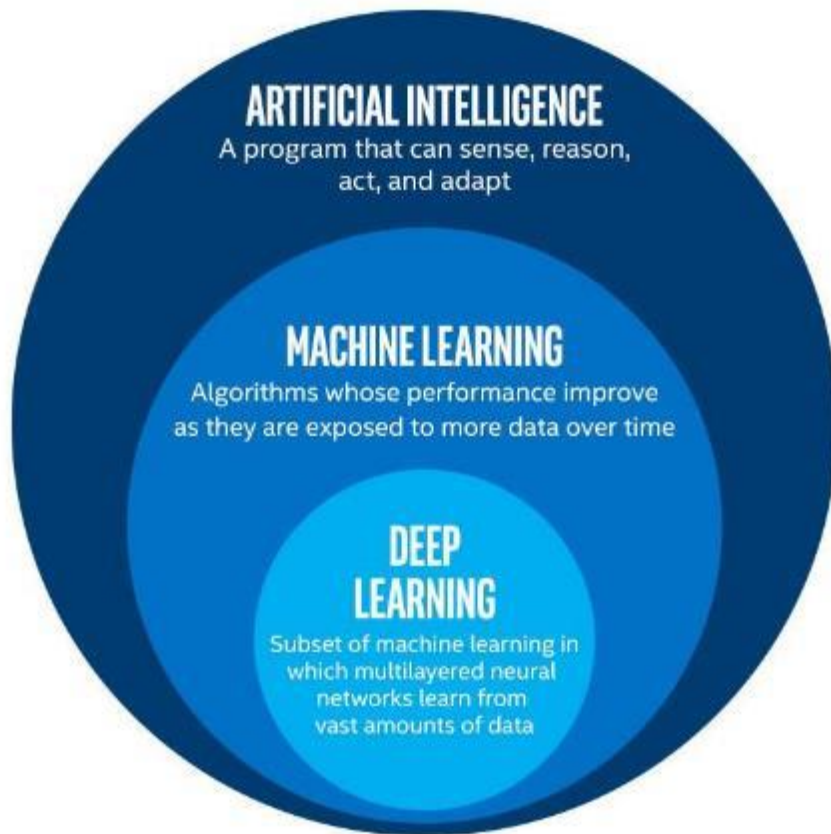


Dog and cat recognition

Deep Learning

“Machine learning that involves using very complicated models called “deep neural networks”.”
(Intel)

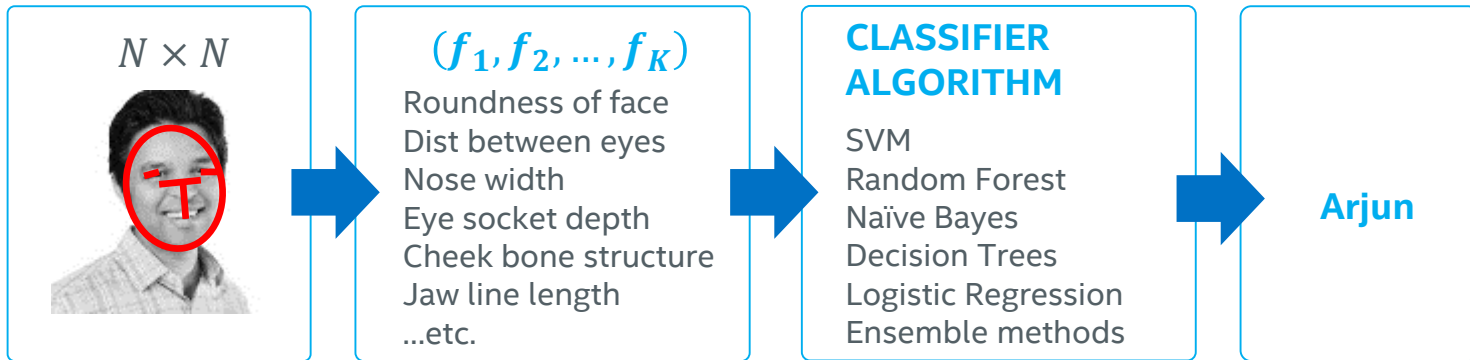
Models determine best representation of original data; in classic machine learning, humans must do this.



MACHINE VS. DEEP LEARNING


MACHINE LEARNING

How do you engineer the best features?



DEEP LEARNING


How do you guide the model to find the best features?




DEEP LEARNING BASICS



TRAINING




INFERENCE



DID YOU KNOW?

Training with a large data set AND deep (many layered) neural network often leads to the highest accuracy inference





MACHINE LEARNING

Machines Learn in Two Ways: Supervised Learning & Unsupervised Learning

Supervised Learning

We train the model. We feed the model with correct answers.
Model Learns and finally predicts.

We feed the model with “ground truth”.

Unsupervised Learning

Data is given to the model. Right answers are not provided to the model. The model makes sense of the data given to it.

Can teach you something you were probably not aware of in the given dataset.

Types of Supervised and Unsupervised learning

SUPERVISED

CLASSIFICATION

REGRESSION

UNSUPERVISED


CLUSTERING

RECOMMENDATION


EXAMPLES OF SUPERVISED LEARNING - CLASSIFICATION

Predict a **label** for an entity with a given set of features.

PREDICTION



SENTIMENT ANALYSIS



Supervised Learning Overview

Training: Train a model with known data.



Inference: Feed unseen data into trained model to make predictions.



Which Model?

Some considerations when choosing are:

- Time needed for training
- Speed in making predictions
- Amount of data needed
- Type of data
- Problem complexity
- Ability to solve a complex problem
- Tendency to overcomplicate a simple one

Evaluation Metric


There are many metrics available* to measure performance, such as:

- **Accuracy:** how well predictions match true values.
- **Mean Squared Error:** average square distance between prediction and true value.

$$\min_{\beta_0, \beta_1} \frac{1}{m} \sum_{i=1}^m \left((\beta_0 + \beta_1 x_{obs}^{(i)}) - y_{obs}^{(i)} \right)^2$$



Accuracy target




**The wrong metric can be misleading or not capture the real problem.*

EXAMPLE OF UNSUPERVISED LEARNING - CLUSTERING

Group entities with similar features

MARKET SEGMENTATION








DEEP LEARNING

DEEP LEARNING BREAKTHROUGHS

Machines able to meet or exceed human image & speech recognition



e.g.




Source: ILSVRC ImageNet winning entry classification error rate each year 2010-2016 (Left), <https://www.microsoft.com/en-us/research/blog/microsoft-researchers-achieve-new-conversational-speech-recognition-milestone/> (Right)

Classification and Detection

Detect and label the image

- Person
- Motor Bike



Semantic Segmentation

Label every pixel



<https://people.eecs.berkeley.edu/~jhoffman/talks/lsta-baylearn2014.pdf>

Natural Language Object Retrieval

a scene with three people query='man far right'



query='man far right'



query='left guy'

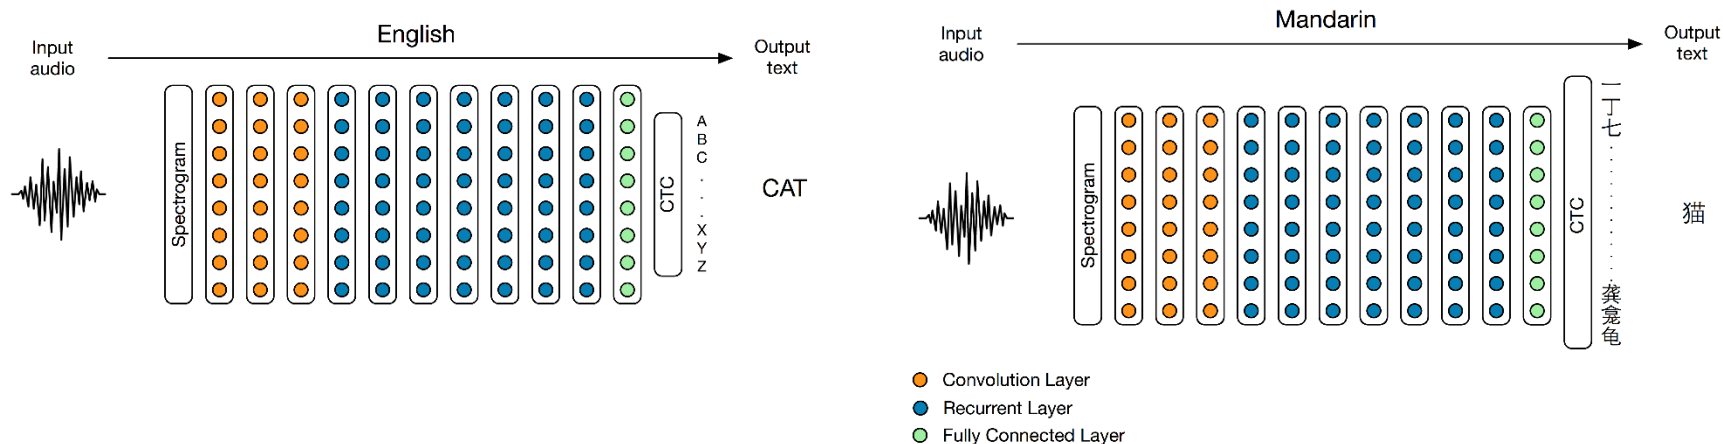


query='cyclist'



<http://arxiv.org/pdf/1511.04164v3.pdf>

Speech Recognition and Language Translation



The same architecture is used for English and Mandarin Chinese speech recognition

<http://svail.github.io/mandarin/>




NEURAL NETWORKS CONNECTIVITY

Fully Connected Network

More complicated problems can be solved by connecting multiple neurons together and using more complicated activation functions.



- Organized into layers of neurons.
- Each neuron is connected to every neuron in the previous layer.
- Each layer transforms the output of the previous layer and then passes it on to the next.
- Every connection has a separate weight



Convolutional Neural Network

Convolutional neural networks reduce the required computation and are good for detecting features.

- Each neuron is connected to a small set of nearby neurons in the previous layer
- The same set of weights are used for each neuron
- Ideal for spatial feature recognition, Ex: Image recognition
- Cheaper on resources due to fewer connections



<http://svail.github.io/mandarin/>



SUMMARY

Summary

- AI on PC will impact and transform many segments in society.
- AI on PC capabilities will continue to grow with ever growing increasing amounts of data available for training along with faster computing power, and better algorithms.
- The differences between machine learning and deep learning are:
 - Machine learning requires humans to engineer the features and the algorithms that will learn either in supervised or unsupervised modes.
 - Deep learning where the algorithms represent a variety of connection between computational nodes – the so called neural-networks.
- The differences between training and inference are:
 - Training refers to the process where data is processed by an algorithm to produce a model that is consistent with the features of the given data.
 - Inference refers to the process where a trained model when presented with new data can make predictions.

