



ARTIFICIAL INTELLIGENCE 501

Lesson 1

An Introduction to AI and Its History

Legal Disclaimers

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

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).

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel and the Intel logo 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.

Learning Objectives

You will be able to:

- Define “Artificial Intelligence” (AI), “Machine Learning” (ML), and “Deep Learning” (DL)
- Explain how DL helps solve classical ML limitations.
- Explain key historical developments, and the “hype-AI winter cycle.”
- Differentiate modern AI from prior AI.
- Relate sample applications of AI.



AI Breakthroughs

Image classification



“Dog”

“Cat”

As of 2015, computers can be trained to perform better on this task than humans.

Machine translation

“I am a student”



“Je suis étudiant”


As of 2016, we have achieved near-human performance using the latest AI techniques.

AI Is The New Electricity

“About 100 years ago, electricity transformed every major industry. AI has advanced to the point where it has the power to transform...every major sector in coming years.”

-Andrew Ng, Stanford University

*Projected Revenue (in billions USD)
Generated from AI, 2016-2020 (IDC)*






DEFINITIONS

Definitions

- Artificial Intelligence
- Machine Learning
- Deep Learning



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)

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.



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

- 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 of a cat or 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.



Deep Learning Example

Classic Machine Learning

Step 1: Determine features.
Step 2: Feed them through model.



Feature Detection

Machine Learning Classifier Algorithm


"Arjun"

Deep Learning

Steps 1 and 2 are combined into 1 step.



Neural Network




"Arjun"



HISTORY


History of AI

AI has experienced several hype cycles, where it has oscillated between periods of excitement and disappointment.



1950s: Early AI

- 1950: Alan Turing developed the Turing test to test a machine's ability to exhibit intelligent behavior.
- 1956: Artificial Intelligence was accepted as a field at the Dartmouth Conference.
- 1957: Frank Rosenblatt invented the perceptron algorithm. This was the precursor to modern neural networks.
- 1959: Arthur Samuel published an algorithm for a checkers program using machine learning.



The First “AI Winter”

- 1966: ALPAC committee evaluated AI techniques for machine translation and determined there was little yield from the investment.
- 1969: Marvin Minsky published a book on the limitations of the Perceptron algorithm which slowed research in neural networks.
- 1973: The Lighthill report highlights AI’s failure to live up to promises.
- The two reports led to cuts in government funding for AI research leading to the first “AI Winter.”



John R. Pierce, head of ALPAC

1980's AI Boom

- Expert Systems - systems with programmed rules designed to mimic human experts.
- Ran on mainframe computers with specialized programming languages (e.g. LISP).
- Were the first widely-used AI technology, with two-thirds of "Fortune 500" companies using them at their peak.
- 1986: The "Backpropagation" algorithm is able to train multi-layer perceptrons leading to new successes and interest in neural network research.



Early expert systems machine

Another AI Winter (late 1980's – early 1990s)

- Expert systems' progress on solving business problems slowed.
- Expert systems began to be melded into software suites of general business applications (e.g. SAP, Oracle) that could run on PCs instead of mainframes.
- Neural networks didn't scale to large problems.
- Interest in AI in business declined.

Late 1990's to early 2000's: Classical Machine Learning

- Advancements in the SVM algorithm led to it becoming the machine learning method of choice.
- AI solutions had successes in speech recognition, medical diagnosis, robotics, and many other areas.
- AI algorithms were integrated into larger systems and became useful throughout industry.
- The Deep Blue chess system beat world chess champion Garry Kasparov.
- Google search engine launched using artificial intelligence technology.



IBM supercomputer

2006: Rise of Deep Learning

- 2006: Geoffrey Hinton publishes a paper on unsupervised pre-training that allowed deeper neural networks to be trained.
- Neural networks are rebranded to deep learning.
- 2009: The ImageNet database of human-tagged images is presented at the CVPR conference.
- 2010: Algorithms compete on several visual recognition tasks at the first ImageNet competition.

IMGENET



MODERN AI

Deep Learning Breakthroughs (2012 – Present)

- In 2012, deep learning beats previous benchmark on the ImageNet competition.
- In 2013, deep learning is used to understand “conceptual meaning” of words.
- In 2014, similar breakthroughs appeared in language translation.
- These have led to advancements in Web Search, Document Search, Document Summarization, and Machine Translation.



Google Translate

Deep Learning Breakthroughs (2012 – Present)

- In 2014, computer vision algorithm can describe photos.
- In 2015, Deep learning platform TensorFlow* is developed.
- In 2016, DeepMind* AlphaGo, developed by Aja Huang, beats Go master Lee Se-dol.



Modern AI (2012 – Present): Deep Learning Impact

Computer vision



Self-driving cars:
object detection




Healthcare:
improved diagnosis

Natural language



Communication:
language translation

How Is This Era of AI Different?



Other Modern AI Factors

- Continued expansion of open source AI, especially in Python*, aiding machine learning and big data ecosystems.
- Leading deep learning libraries *open sourced*, allowing further adoption by industry.
- Open sourcing of large datasets of millions of labeled images, text datasets such as Wikipedia has also driven breakthroughs.



Transformative Changes



Health

Enhanced
Diagnostics
Drug Discovery
Patient Care
Research
Sensory Aids



Industrial

Factory
Automation
Predictive
Maintenance
Precision
Agriculture
Field
Automation

Source: Intel forecast

Transformative Changes



Finance

Algorithmic Trading
Fraud Detection
Research
Personal Finance
Risk Mitigation



Energy

Oil & Gas Exploration
Smart Grid
Operational Improvement
Conservation

Source: Intel forecast

Transformative Changes



Government

- Defense
- Data Insights
- Safety & Security
- Engagement
- Smarter Cities



Transport

- Autonomous Cars
- Automated Trucking
- Aerospace
- Shipping
- Search & Rescue

Source: Intel forecast

Transformative Changes



Other

Advertising

Education

Gaming

Professional &
IT Services

Telco/Media

Sports


Source: Intel forecast



APPLICATIONS

AI Omnipresence In Transportation

Navigation



Google & Waze find the fastest route, by processing traffic data.

Ride sharing



Uber & Lyft predict real-time demand using AI techniques, machine learning, deep learning.

AI Omnipresence In Social Media

Audience



Facebook & Twitter use AI to decide what content to present in their feeds to different audiences.

Content



Image recognition and sentiment analysis to ensure that content of the appropriate “mood” is being served.

AI Omnipresence In Daily Life

Natural language



We carry around powerful natural language processing algorithms in our phones/computers.

Object detection




Cameras like Amazon DeepLens* or Google Clips* use object detection to determine when to take a photo.

Latest Developments: Computer Vision




mite	container ship	motor scooter	leopard
mite	container ship	motor scooter	leopard
black widow	lifeboat	go-kart	jaguar
cockroach	amphibian	moped	cheetah
tick	fireboat	bumper car	snow leopard
starfish	drilling platform	golfcart	Egyptian cat

Deep Learning
“proven” to work for
image classification.



Models outperform
humans on image
classification.



Object detection
models beat previous
benchmarks.


2012

2015

2016

Application Area: Abandoned Baggage Detection

- We can automatically detect when baggage has been left unattended, potentially saving lives.
- This system relies on the breakthroughs we discussed:
 - Cutting edge object detection.
 - Fast hardware on which to train the model (Intel® Xeon® processors in this case).




Abandoned baggage

Learning Objectives Recap

In this session, we worked to:

- Define “Artificial Intelligence” (AI),
“Machine Learning” (ML), and “Deep Learning” (DL).
- Explain how DL helps solve classical ML limitations.
- Explain key historical developments,
and the “hype-AI winter cycle.”
- Differentiate modern AI from prior AI.
- Relate sample applications of AI.





Sources for information and images used in this presentation

Ai face image: ✓ <https://pixabay.com/en/girl-woman-face-eyes-close-up-320262/>

Dog image: ✓ <https://www.pexels.com/photo/adorable-animal-breed-canine-356378/>

Cat image: ✓ <https://www.pexels.com/photo/adorable-animal-baby-blur-177809/>

Brain image: ✓ <http://www.publicdomainpictures.net/view-image.php?image=130359&picture=human-brain>

Email icon: ✓ <https://pixabay.com/en/mail-message-email-send-message-1454731/>

Credit card: ✓ <https://www.pexels.com/photo/atm-atm-card-card-credit-card-371066/>

Dog and cat image: ✓ <https://pixabay.com/en/cat-dog-animals-pets-garden-71494/>

John Pierce: ✓ https://commons.wikimedia.org/wiki/File:John_Robinson_Pierce.jpg

Sources for information and images used in this presentation

Mainframe Computer: ✓

https://en.wikipedia.org/wiki/IBM_System/360_Model_40#/media/File:IBM_System_360_at_USDA.jpg

Car: ✓ <https://upload.wikimedia.org/wikipedia/commons/thumb/3/3e/Stanley2.JPG/1200px-Stanley2.JPG>

Spotify/netflix: ✓ <https://www.pexels.com/photo/spotify-netflix-mockup-364676/>

Googlecar: ✓

[https://commons.wikimedia.org/wiki/File:Secretary_Kerry_Views_the_Computers_Inside_One_of_Google%27s_Self-Driving_Cars_at_the_2016_Global_Entrepreneurship_Summit%27s_Innovation_Marketplace_at_Stanford_University_\(27786624121\).jpg](https://commons.wikimedia.org/wiki/File:Secretary_Kerry_Views_the_Computers_Inside_One_of_Google%27s_Self-Driving_Cars_at_the_2016_Global_Entrepreneurship_Summit%27s_Innovation_Marketplace_at_Stanford_University_(27786624121).jpg)

Emoji: ✓ <https://pixabay.com/en/emoji-emoticon-smilies-icon-faces-2074153/>

Sources for information and images used in this presentation

Stethoscope: ✓ <https://pixnio.com/science/medical-science/stethoscope-mobile-phone-diagnosis-pulse-cardiology-medicine-healthcare>

Translation keys: ✓ <https://www.pexels.com/photo/interpretation-transcription-translation-569313/>

Googletranslate: ✓ https://commons.wikimedia.org/wiki/File:Google_Translate_logo.svg

Big data: ✓ <https://pixabay.com/en/big-data-data-world-cloud-1667184/>

Google pin: ✓ <https://pixabay.com/en/navigation-gps-location-google-2049643/>

Suv: ✓ <http://maxpixel.freegreatpicture.com/Land-Truck-Range-Vehicle-Range-Rover-Car-4-Rover-2015663>

Social media: ✓ <https://pixabay.com/en/social-media-manager-online-woman-1966010/>

Girlphone: ✓ <https://pixnio.com/people/female-women/digital-camera-mobile-phone-fashion-girl-smartphone-woman>

Phonehand: ✓ <https://www.pexels.com/photo/internet-mobile-phone-phone-screen-270661/>