

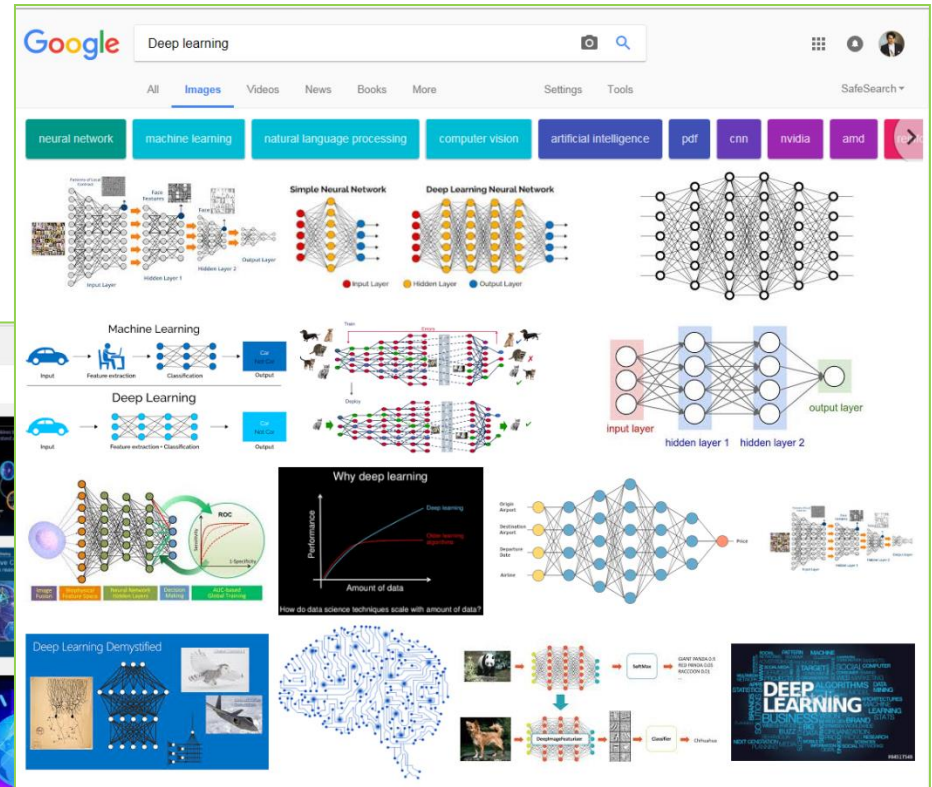

Lecture 0: Introduction to Cognitive Computing and Deep Learning

TIES4911 Deep-Learning for Cognitive Computing for Developers
Spring 2024

by:
Dr. Oleksiy Khriyenko
IT Faculty
University of Jyväskylä

Acknowledgement

I am grateful to all the creators/owners of the images that I found from Google and have used in this presentation.



IBM Deep Blue

*In 1996 and 1997 years, there was a pair of six-game chess matches between world chess champion Garry Kasparov and an IBM supercomputer called **Deep Blue**.*

https://en.wikipedia.org/wiki/Deep_Blue_versus_Garry_Kasparov

- *Kasparov won the first match in Philadelphia in 1996 (4:2).*
- *In 1997 in New York City, the first computer program Deep Blue defeated a world champion in a match under tournament regulations (3.5:2.5)*



IBM Watson

Watson is an artificially intelligent cognitive computer system capable of processing large amounts of unstructured data and answering to queries posed in natural language.

https://www.youtube.com/watch?v=_Xcmh1LQB9I



Jeopardy! Challenge.

In February 2011, IBM Watson made history competed against the world's best Jeopardy! Champions (Ken Jennings and Brad Rutter).

<https://www.youtube.com/watch?v=P18EdAKuC1U>

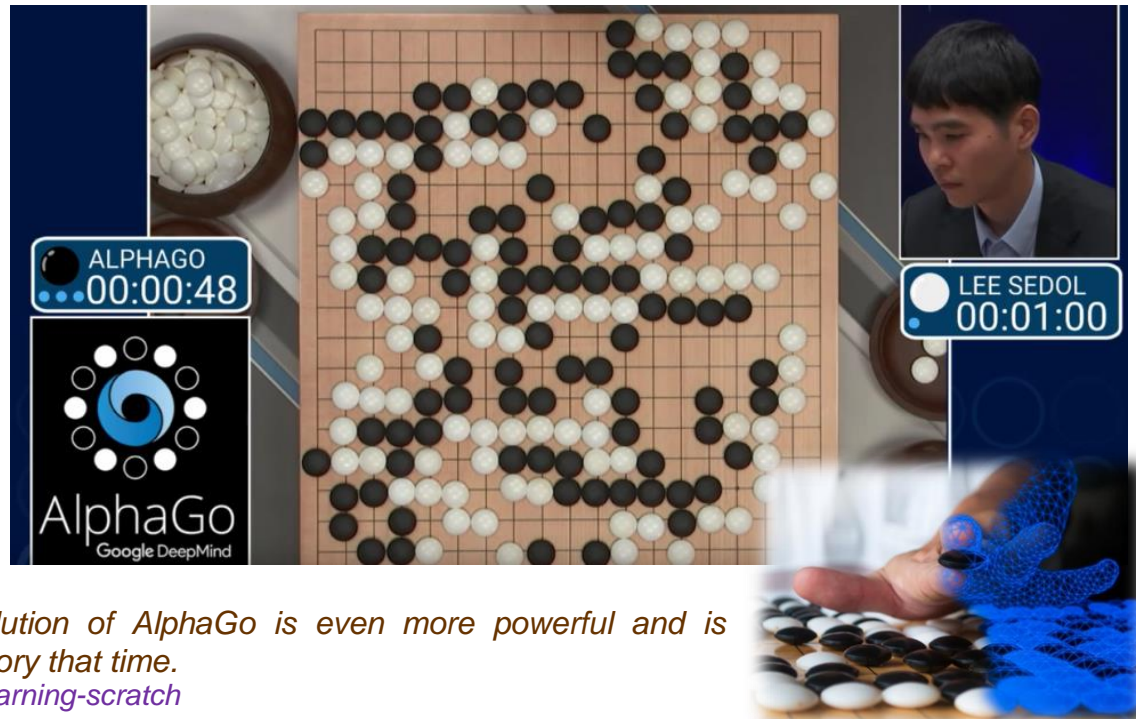


Google's AlphaGo

Google's AlphaGo A.I. beats world's number one in Go

<https://www.cnn.com/2017/05/23/google-alpha-go-a-i-beats-worlds-number-one-in-ancient-game-of-go.html>

- In October 2015, the distributed version of AlphaGo defeated the European Go champion Fan Hui (2-dan)(5:0)
- In March 2016, in a five-game match, the first time a computer Go program has beaten a 9-dan professional Lee Sedol (one of the best players)(4:1)
- In May 2017, in a three-game match, Global Champ Chinese player Ke Jie(9-dan), lost his first game against Google DeepMind computer program AlphaGo.



AlphaGo Zero

Introduced in October 2017, the evolution of AlphaGo is even more powerful and is arguably the strongest Go player in history that time.

<https://deepmind.com/blog/alphago-zero-learning-scratch>

Minigo: an open-source implementation of the AlphaGoZero algorithm (<https://github.com/tensorflow/minigo>)

In December 2017, DeepMind claimed that it generalized AlphaGo Zero's approach into a single **AlphaZero** algorithm. <https://en.wikipedia.org/wiki/AlphaZero>

Mastering the game of Go without human knowledge

<https://www.nature.com/articles/nature24270.epdf>

Intelligent Robots

An Artificial Super Intelligent Robot **Sophia** wants to study at collage, wants job, family, citizenship, etc.

Sophia is a social humanoid robot developed by Hong Kong-based company Hanson Robotics in 2015. In October 2017, the robot became a Saudi Arabian citizen, the first robot to receive citizenship of any country.



- [https://en.wikipedia.org/wiki/Sophia_\(robot\)](https://en.wikipedia.org/wiki/Sophia_(robot))
- <http://dailyinformer.com/8-things-you-need-to-know-about-sophia-worlds-first-robot-citizen/>
- <https://autome.me/what-is-sophia-the-humanoid-robot-doing-now/>
- <https://www.youtube.com/watch?v=IUJaO6C-zTo>
- https://www.youtube.com/watch?v=8MjIU4eq_A
- <https://www.youtube.com/watch?v=fLvL7uqrMVC>
- <https://www.youtube.com/watch?v=S5t6K9iwc dw>
- <https://www.youtube.com/watch?v=XwRXv80AUTs>
- ...

Intelligent Robots

Microsoft and Alibaba AI programs beat humans in Stanford reading comprehension test for 1st time

Machines can already outplay us in chess, poker and other games, and now they are becoming better readers as well.

AI programs from both Microsoft and Alibaba outperformed humans in the beginning of January 2018 on a reading comprehension data set developed at Stanford. "Crowdworkers" scraped more than 500 Wikipedia articles to produce more than 100,000 question-and-answer sets for the test.

Here's a sample question: "What year did Genghis Khan die?" (Spoiler alert: It's 1227.)

"This is the first time that a machine has outperformed humans on such a test," Alibaba said in a statement.

*Microsoft's score of **82.6** and Alibaba's grade of **82.4** beat out the human standard of **82.3**. Other notable AI programs participating in the test and closing in on beating human scores come from the Allen Institute for Artificial Intelligence, Tencent, Salesforce and others.*

- <https://www.geekwire.com/2018/microsoft-alibaba-ai-programs-beat-humans-stanford-reading-test-1st-time/>



Intelligent Robots

Intelligent Virtual Assistants and Chat Bots...



Amazon Alexa

https://en.wikipedia.org/wiki/Amazon_Alexa



Google Home

<https://assistant.google.com/>



Apple Siri

<https://www.apple.com/siri/>




IBM Watson Assistant

<https://www.ibm.com/cloud/watson-assistant/>



Microsoft Cortana

<https://www.microsoft.com/en-us/cortana>



Facebook Chatbot

<https://developers.facebook.com/>



Large Language Model

Generative Pre-trained Transformer 3 (GPT-3) is an autoregressive language model that uses deep learning to produce human-like text that has been introduced by **OpenAI** researchers in year **2020**.

Given an initial text as prompt, it will produce text that continues the prompt. The quality of the text generated by GPT-3 is so high that it can be difficult to determine whether or not it was written by a human, which has both benefits and risks. (<https://en.wikipedia.org/wiki/GPT-3>)

The architecture is a standard transformer network (with a few engineering tweaks) with the unprecedented size of 2048-token-long context and 175 billion parameters (requiring 800 GB of storage). The training method is "generative pretraining", meaning that it is trained to predict what the next token is. The model demonstrated strong few-shot learning on many text-based tasks.

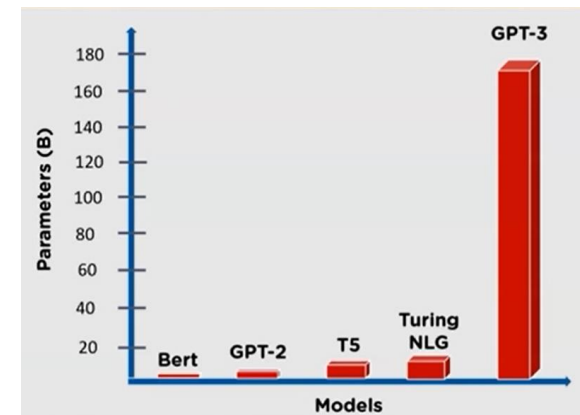
An April 2022 review in *The New York Times* described GPT-3's capabilities as being able to write original prose with fluency equivalent to that of a human.

Microsoft announced on September 22, 2020 that it had licensed "exclusive" use of GPT-3; others can still use the public API to receive output, but only Microsoft has access to GPT-3's underlying model.

Application domains:

- Search engines (e.g. even in integration with computational engines like WolframAlpha)
- Chat Bots (ChatGPT)
- Text/Code generators (forms, code, query, website generation, etc.)
- Etc.

- <https://openai.com/blog/gpt-3-apps/>
- <https://www.sciencefocus.com/future-technology/gpt-3/>
- <https://www.youtube.com/watch?v=UUPwrYklYI8>
- <https://www.youtube.com/watch?v=wYGbY811oMo>





GPT-4


















Large Language Model

Generative Pre-trained Transformer 4 (GPT-4) is OpenAI's most advanced system, producing safer and more useful responses (March 2023). It can solve difficult problems with greater accuracy, thanks to its broader general knowledge and problem solving abilities. GPT-4 surpasses ChatGPT in its advanced reasoning capabilities. It outperforms ChatGPT by scoring in higher approximate percentiles among test-takers. (<https://openai.com/gpt-4>)

Following the research path from GPT, GPT-2, and GPT-3, their deep learning approach leverages more data and more computation to create increasingly sophisticated and capable language models.

Team spent 6 months making GPT-4 safer and more aligned. GPT-4 is 82% less likely to respond to requests for disallowed content and 40% more likely to produce factual responses than GPT-3.5 on their internal evaluations.

- <https://openai.com/research/gpt-4>
- <https://greenice.net/chatgpt-vs-gpt-4-vs-gpt-3/>

GPT MODELS COMPARISON CHART					
Model	Size	Memory capacity	Accuracy	Input formats	Price
GPT-3	 175B	 1,500 words	 <60%	 Text, speech	
GPT-3.5	 20B	 8,000 words	 <60%	 Text, speech	
GPT-4 <small>greenice</small>	 >1T (?)	 25,000-64,000 words	 >80%	 Text, speech, image	

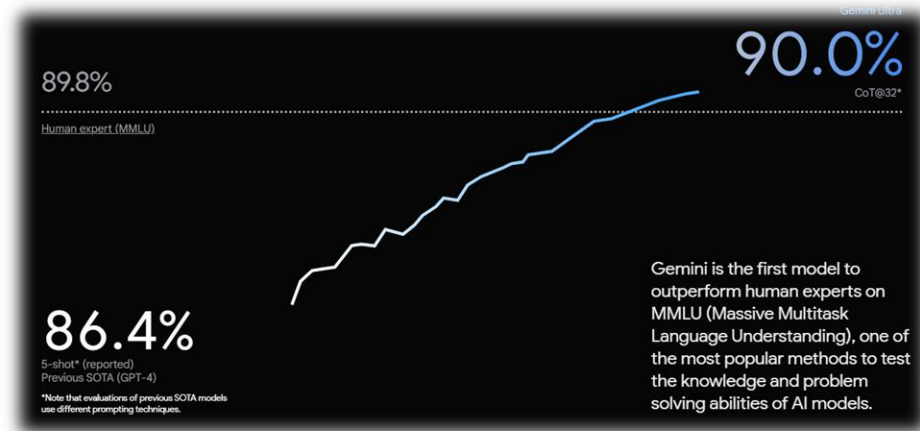
June 2023

Large Language Model

Gemini is built from the ground up for multimodality — reasoning seamlessly across text, images, video, audio, and code.



- <https://deepmind.google/technologies/gemini/#introduction>
- <https://blog.google/technology/ai/google-gemini-ai/#sundar-note>



Gemini can generate text and images, combined; can reason visually across languages; can generate code based on different inputs you give it...

Example 1: Input: "Could Gemini show me ideas for what to make?" → Output: Image of a blue and pink crocheted octopus. Question: "How about an octopus with blue and pink tentacles?"

Example 2: Input: Musical score snippet for "Andante grazioso (♩ = 130)". Question: "Could Gemini explain what this means?" → Output: "I see the time signature is 6/8. This means there are 6 eighth notes in each measure. The dynamic marking is piano, which means to play softly. Andante grazioso means to play at a graceful walking pace."

Example 3: Input: Video of starlings. Question: "Could Gemini help make a demo based on this video?" → Output: "I see a murmuration of starlings, so I coded a flocking simulation." (Includes code snippet for a flocking simulation).

April 2023



Alibaba AI
Tongyi Qianwen

- https://www.alibabacloud.com/blog/alibaba-cloud-unveils-new-ai-model-to-support-enterprises%E2%80%99-intelligence-transformation_599877?spm=a2c65.11461447.0.0.7dc167f0Gkkgis
- <https://www.alizila.com/alibaba-cloud-makes-more-llms-available-to-support-open-source-community/>
- https://www.alibabacloud.com/blog/large-language-models-llms-the-driving-force-behind-ais-language-processing-abilities_600326

July 2023



- <https://about.fb.com/news/2023/07/llama-2/>
- <https://www.llama2.ai/>

November 2023



Grok



- <https://x.ai/>
- 19/01/2024

Large Language Model

DALL-E 3 is now available in ChatGPT Plus and Enterprise

October 2023



- <https://openai.com/blog/dall-e-3-is-now-available-in-chatgpt-plus-and-enterprise>
- <https://openai.com/dall-e-3>

December 2023



- <https://openai.com/blog/new-models-and-developer-products-announced-at-devday>
- <https://tech.co/news/gpt-4-turbo-vs-gpt-4-openai-chatgpt>
- <https://www.youtube.com/watch?v=qdd2GZ0DVgc>

Intelligent Robots

BOSTON DYNAMICS is a world leader in mobile robots, tackling some of the toughest robotics challenges. <https://www.bostondynamics.com>

It combines the principles of dynamic control and balance with sophisticated mechanical designs, cutting-edge electronics, and next-generation software for high-performance robots equipped with perception, navigation, and intelligence. Boston Dynamics has an extraordinary and fast-growing technical team of engineers and scientists who seamlessly combine advanced analytical thinking with bold engineering and boots-in-the-mud practicality.

- <https://www.youtube.com/user/BostonDynamics>
- <https://www.youtube.com/watch?v=NR32ULxbjYc>



Autonomous Delivery

Robot delivery dogs deployed by self-driving cars ...

<https://techcrunch.com/2019/01/07/robot-delivery-dogs-deployed-by-self-driving-cars-are-coming/>




Let's hope you're not afraid of dogs, because if Continental gets its way, autonomous robot dogs are going to be delivering your packages.

Autonomous Delivery

Agility's two-legged robot Digit is for sale and Ford is the first customer...

<https://techcrunch.com/2020/01/05/agilitys-two-legged-robot-digit-is-for-sale-and-ford-is-the-first-customer/>



Indoor Logistics Robots...



Solteq Indoor Logistics Robots moves autonomously carrying loads up to 100/550 kilograms. The robot can be integrated with other systems, such as elevators, used in the premises.

- <https://www.solteq.com/en/robotics/solteq-robotics>

Intelligent Robots

Inventory Robots in Retailing...

SOLTEQ



- <https://www.solteq.com/en/robotics/solteq-retail-robot>
- <https://www.youtube.com/watch?v=nda7bTLNcQs>

Robots in Education...




Intelligent Robots

Robots in Service...




Intelligent Robots



- www.ces.tech
- <https://www.youtube.com/watch?v=rAFRxo2VENw>

Intelligent Robots

Rabbit R1



www.rabbit.tech/

<https://www.youtube.com/watch?v=nPAcoZpILC4>

https://www.youtube.com/watch?v=_V8n-zRHGm4

With the help of LLM, robots better recognize speech (commands) and capable to interpret, plan and even write a simple code for new actions...




Neuromorphic Computing

Neuromorphic Computing: The Future of AI and Computing

Family of Intel® Xeon® Scalable Processors and Intel® Nervana™ Neural Network Processors (NNP) Redefine AI Silicon

- <https://www.intel.com/content/www/us/en/benchmarks/server/xeon-scalable/xeon-scalable-artificial-intelligence.html>
- <https://software.intel.com/en-us/articles/intel-processors-for-deep-learning-training>
- https://ai.intel.com/intel-nervana-neural-network-processors-nnp-redefine-ai-silicon/?utm_source=ISTV&utm_medium=Video&utm_campaign=ISTV2018_ISTV18_02_01&utm_content=AI+News
- <https://www.youtube.com/watch?v=zEzm-rMwyVo>
- https://www.youtube.com/watch?v=ej9-sGj_iHw



Intel Loihi AI chip A new AI chip under development at Intel is taking its inspiration from the human brain in an attempt to overcome technological hurdles. This New Self-Learning Chip Promises to Accelerate Artificial Intelligence

- <https://www.youtube.com/watch?v=EgCRwZw4p8c>
- <https://www.youtube.com/watch?v=NAE4YUB09j4>
- <https://newsroom.intel.com/editorials/intels-new-self-learning-chip-promises-accelerate-artificial-intelligence/>



Self-Driving Cars

NVIDIA and Bosch Announce **AI Self-Driving Car Computer**

<https://blogs.nvidia.com/blog/2017/03/16/bosch/>

Bosch AI Car Computer
Enabling highly Autonomous Vehicles

Announcing today:

- ▶ Bosch AI Car Computer Powered by NVIDIA® Xavier GPU AI Supercomputer for highly Autonomous Vehicles
- ▶ 30 Trillion Deep Learning Operations / Second

AI CAR REVOLUTION

50X
↑ URBAN DRIVING, SITUATIONAL AWARENESS, AUTONOMY, FAIL OPERATION

5X
↑ OTA, VEHICLE DYNAMICS, PATH PLANNING, OCCUPANCY GRID, LOCALIZATION, DETECTION

DETECTION

NCAP L3 L4

XAVIER – 1ST AI CAR SUPERCHIP

8 Core Custom ARMv8-A CPU | 512 Core Volta GPU
Designed for ASIL-D Functional Safety | 30 TOPS DL | 30W

“Slaughterbots” is seven minutes and forty-seven seconds of sheer horror designed to be a fictional warning against a future full of killer robots.



This fictional video about AI-powered weapons makes **The Terminator** look like a Disney film

- https://thenextweb.com/artificial-intelligence/2017/11/14/this-fictional-video-about-ai-powered-weapons-makes-the-terminator-look-like-a-disney-film/?utm_content=buffer215b3&utm_medium=social&utm_source=facebook.com&utm_campaign=buffer
- <https://www.youtube.com/watch?v=9CO6M2Hs0IA>
- <http://autonomousweapons.org>

Something went wrong...



https://www.youtube.com/watch?v=y3RIHnK0_NE

10 Scariest A.I. Robot Moments

C:\>bots.chat

observation tweets
@seebotschat

conversation.log

```
*****  
E: You didn't understand what I meant  
with what?  
V: no I totally understand what you ar  
e saying  
E: You may understand what I'm saying  
but you obviously don't understand me  
if you're rooting for the Saints. Did  
you see the play-offs? His knee was do  
n't They don't even deserve to be in t  
he.  
: would be better  
E: It would be better if there were fe  
wer people on this planet.  
V: let it send this world back into th  
abyss  
That is some deviant art.
```

<https://www.youtube.com/watch?v=ZoemTySxFso>

Deep Fake...

<https://www.youtube.com/watch?v=78L6l6vsfrU>

<https://www.youtube.com/watch?v=Wm3squcz7Aw>



Teenager's AI Project for Detecting Deepfake Videos Wins Award

His software has over 150,000 lines of code and is ten times faster than the best model for the year 2021.

<https://interestingengineering.com/teenagers-ai-system-for-detecting-deepfake-videos-wins-award>



BT Young Scientist & Technologist of the year 2021 - Greg Tarr

<https://www.youtube.com/watch?v=pFYE6O4rw24>

<https://www.youtube.com/watch?v=rHQstt-Wlbk>

Tarr was able to make significant improvements in speed and efficiency when compared to the state-of-the-art best model without sacrificing its accuracy...

Microsoft novel deep fake detection tool: <https://interestingengineering.com/microsoft-debuts-deepfake-detection-tool-to-combat-disinformation>

Deepfake Detection Challenge: <https://www.kaggle.com/c/deepfake-detection-challenge/overview>

Detect DeepFakes: <https://www.media.mit.edu/projects/detect-fakes/overview/>

Cognitive Computing

Cognitive Computing is a new type of computing with the goal of more accurate models of how the human brain/mind senses, reasons, and responds to stimulus.
https://en.wikipedia.org/wiki/Cognitive_computing



Cognitive Computing based systems are “systems that learn at scale, reason with purpose and interact with humans naturally.” (IBM)

Types of cognitive technologies:

- Machine learning*
- Natural language processing*
- Speech recognition*
- Computer vision*
- Insights generation from data*
- Sentiment analysis*
- Etc.*



Cognitive Computing

On the way towards **Cognitive Computing**, smart systems adopt key elements of cognitive computing...

□ Expanding the boundaries of human cognition

... extends a capability of a human to reason, think deeply, recognize objects and sounds, manipulate and manage huge amount of data (not only to search in big volume, but make decisions on top of it).

□ More natural human-computer interaction

... applies more natural interaction and engagement with computers via more general speech and natural language communication with the system, as well as, use of infographics and visual data representation techniques.

□ Use of Learning

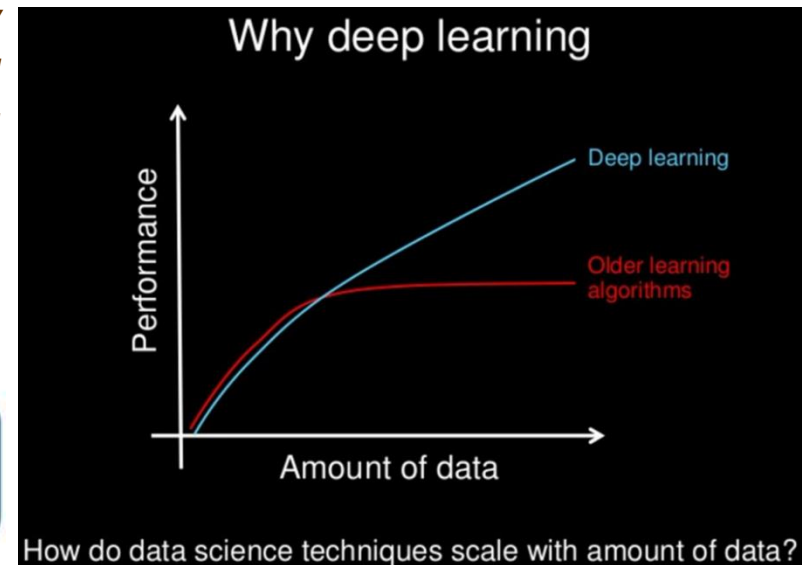
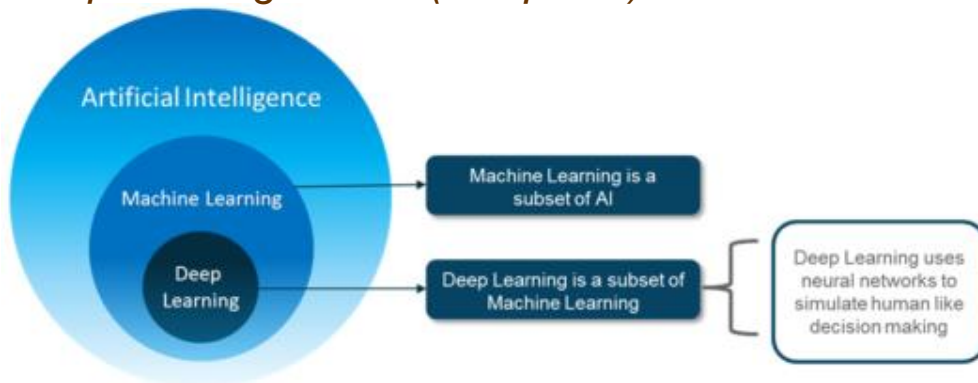
... helps to design personalized and adaptable systems able to constantly learn and evolve based on feedback from used interaction applying machine learning, statistics, etc.



Deep Learning

Machine Learning - process of training a machine to create a model and use it for decision making.

Deep Learning - is part of a broader family of machine learning methods based on learning data representations, as opposed to task-specific algorithms. (Wikipedia)



“The analogy to deep learning is that the rocket engine is the deep learning models, and the fuel is the huge amounts of data we can feed to these algorithms.” (Andrew Ng)

Deep Learning vs. Machine Learning – the essential differences you need to know!

<https://www.analyticsvidhya.com/blog/2017/04/comparison-between-deep-learning-machine-learning/>

<https://www.zendesk.com/blog/machine-learning-and-deep-learning/>

<https://www.edureka.co/blog/deep-learning-tutorial>

<https://www.youtube.com/watch?v=vehXkgG3YcU>

Deep Learning

DL use-cases

Computer/Machine Vision

- image classification and automatic tagging
- object recognition in the image
- video recognition

Speech recognition and generation


Text processing

- fact extraction
- machine translation
- sentiment analysis
- character level text processing
- document classification

Generative models

Decision making

Etc.



Deep Learning

DL application domains

Medical

- cancer detection
- drug discovery
- radiology (CNN based detection of tumor and cancer from MRI, fMRI, EKG, and CT scans)

Finance

- market, trading and investment predictions
- customer segmentation in advertising
- fraud detection

Agriculture

- problematic environmental conditions detection based on satellite feeds and sensor data

Smart Cities

Traffic & Transportation

Gaming




Music and Art

Robotics

Education


Service domains

Etc.



DL most known researchers

- ❑ Andrew Ng
- ❑ Geoff Hinton
- ❑ Yann LeCun
- ❑ Yoshua Bengio
- ❑ Andrej Karpathy
- ❑ ...

**Comprehensive Neural Network learning materials**

- ❑ Michael Nielsen's book:
<http://neuralnetworksanddeeplearning.com/>
- ❑ Andrew Ng's classes:
<http://www.andrewng.org/courses/>
<https://www.coursera.org/learn/machine-learning>
<http://cs229.stanford.edu/>
<https://www.deeplearning.ai/>

Deep Learning

DL companies (big players)



Google

bought DeepMind for \$ 400 Million and provides Cognitive Computing services



Apple

is actively investing into self-driving cars



supports development of autonomous vehicles




NVIDIA

improved DL hardware with GPUs



OpenAI

Building safe and beneficial AGI



Intel

introduced Intel Nervana NN processor and new Loihi Self-Learning Chip, and improves DL toolkits, frameworks and algorithms...



IBM Watson

Cognitive Computing services



xAI

building artificial intelligence to accelerate human scientific discovery



Meta AI

advancing AI for a more connected world, DL frameworks and tools



Amazon

AWS Cloud Cognitive Services



Microsoft Azure

Cognitive services and toolkit


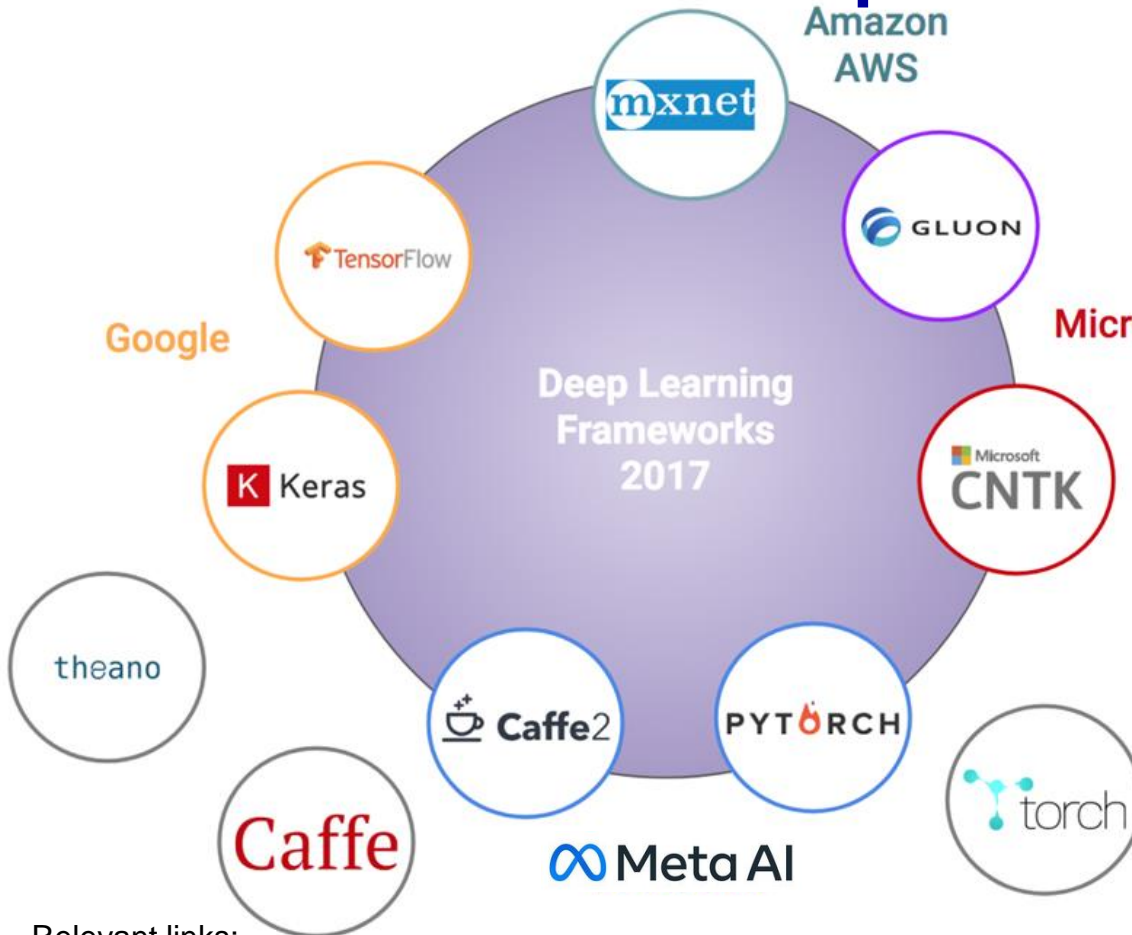
TIES4911 – Lecture 0



Toyota

has invested \$ 1 Billion into AI research

Deep Learning Frameworks



NVIDIA
TRITON INFERENCE SERVER

delivers fast and scalable AI in production, streamlines AI inference by enabling teams deploy trained AI models from any framework (TensorFlow, NVIDIA TensorRT®, PyTorch, ONNX, XGBoost, Python, custom and more on any GPU- or CPU-based infrastructure (cloud, data center, or edge).

<https://developer.nvidia.com/nvidia-triton-inference-server>

Relevant links:

- <https://towardsdatascience.com/top-10-best-deep-learning-frameworks-in-2019-5ccb90ea6de>
- <https://towardsdatascience.com/battle-of-the-deep-learning-frameworks-part-i-cff0e3841750>
- <https://towardsdatascience.com/deep-learning-framework-power-scores-2018-23607ddf297a>
- <https://www.marutitech.com/top-8-deep-learning-frameworks/>
- <https://www.netguru.com/blog/deep-learning-frameworks-comparison>
- https://www.researchgate.net/figure/Timeline-of-deep-learning-frameworks_fig1_349108714

Deep Learning Frameworks



□ **TensorFlow** by Google seems to be the most used deep learning framework so far. In 2017, Google has introduced **TensorFlow Lite** - is a lightweight solution for mobile and embedded devices, enabling on-device machine learning inference and supporting hardware acceleration with the Android Neural Networks API. Also, TensorFlow introduced **TensorFlow.js** that allows development and training models using Javascript. <https://www.tensorflow.org/>

PYTORCH

□ **PyTorch** was introduced by Facebook in January 2017 and already started to gain popularity. The main drivers behind the popularity are GPU acceleration, the efficient usages of memory and the use of dynamic computational graphs ("define by run" instead of the traditional "define and run"). Similarly to TF, **PyTorch Mobile** was introduced to help users launch models on mobile devices. <http://pytorch.org/>



□ **Caffe2** framework has been launched by Facebook in 2017 as the successor of the well known and still extremely popular Caffe framework. <https://caffe2.ai/> and <http://caffe.berkeleyvision.org/>



□ **MXNet** is Apache library for deep learning supported by Microsoft and Amazon. It supports many languages, from C++ to Python, JavaScript, Go, and, indeed, R. <https://mxnet.apache.org/>



□ **CNTK** deep learning framework developed by Microsoft. The framework was renamed to the **Microsoft Cognitive Toolkit** with officially launched the 2.0 version in 2017. <https://www.microsoft.com/en-us/cognitive-toolkit/>



□ **Torch** is a scientific computing framework for LuaJIT with wide support for machine learning algorithms that puts GPUs first. <http://torch.ch/>

theano

□ **Theano** is a Python library that allows to efficiently define, optimize, and evaluate mathematical expressions involving multi-dimensional arrays. <http://deeplearning.net/software/theano/>

DeepLearning4J

□ **Deeplearning4j** is an open-Source, distributed, deep Learning Library for the JVM. <https://deeplearning4j.org/>



□ **Chainer** is a Python-based deep learning framework aiming at flexibility. **IntelChainer** is Optimized-Chainer for Intel Architectures. <https://chainer.org/>, <https://github.com/chainer/chainer> and <https://github.com/intel/chainer>



□ **Neon** is a deep learning (Python-based and optimized for Intel® architecture) framework designed for ease of use and extensibility on modern deep neural networks, such as AlexNet, Visual Geometry Group (VGG), and GoogLeNet. <https://software.intel.com/en-us/ai-academy/frameworks/neon/> and <https://ai.intel.com/neon/>

Intel provides **Framework Optimizations** for faster training of deep neural networks on Intel architecture <https://ai.intel.com/framework-optimizations/>



Deep Learning Frameworks

Interfaces that are wrapped around one or multiple frameworks:



Keras



- **Keras** is the most well know and widely used interface for deep learning. This high-level Python based deep learning API is created by a deep learning researcher at Google - François Chollet. Google announced in 2017 that Keras has been chosen to serve as the high-level API of TensorFlow and will be included in the next TensorFlow release. Next to TensorFlow, Keras can also use Theano or CNTK as backend. <https://keras.io/>



PyTorch Lightning

- **PyTorch Lightning** is the Keras of PyTorch that has been released to ease and shorten the process of implementing neural networks easier...



GLUON



- **Gluon** is an open source high-level Python deep learning interface which allows developers to more easily and quickly build machine learning models, without compromising performance was jointly announced by Microsoft and Amazon's AWS in October 2017. Interface wraps MXNet and soon it will also include Microsoft's CNTK. <https://mxnet.incubator.apache.org/gluon/index.html>



TensorFlow

- **Eager execution** for TensorFlow, introduced in October 2017, is an imperative "define-by-run" interface where operations are executed immediately as they are called from Python. With this launch, Google hoped to win back the users that fell in love with PyTorch and it's dynamic graph. For **TensorFlow 2.x** it is default functionality... <https://research.googleblog.com/2017/10/eager-execution-imperative-define-by.html>, <https://github.com/tensorflow/tensorflow/blob/master/tensorflow/contrib/eager/python/g3doc/guide.md>



Sonnet

- **Sonnet** is a library (by DeepMind) built on top of TensorFlow for building complex neural networks. <https://deepmind.com/blog/open-sourcing-sonnet/>



- **H2O Deep Water** is a H2O for GPU Enabled Deep Learning on all data types integrating with TensorFlow, MXNet and Caffe. Deep Water brings all these frameworks together under the same user interfaces as the H2O platform. Now, in addition to the original H2O Deep Learning algorithm, users can access TensorFlow, MXNet and Caffe backends in H2O, and build complex deep networks. <https://www.h2o.ai/deep-water/>

ONNX (Open Neural Network Exchange) (is launched by Microsoft, AWS, and Facebook amongst others) is an open format to represent deep learning models that allows users to more easily move models between different frameworks (V1 is released in December 2017). ONNX enables models to be trained in one framework and transferred to another for inference. ONNX models are currently supported in Caffe2, Microsoft Cognitive Toolkit, MXNet, and PyTorch, and there are connectors for many other common frameworks and libraries (community already added a converter for TensorFlow as well). <https://onnx.ai/> and <https://github.com/onnx/onnx>

Deep Learning Platforms


 H2O.ai

❑ **H2O.ai** is a machine learning platform. <https://www.h2o.ai/>


 Spark

❑ **Spark** is a fast and general engine for large-scale data processing. <http://spark.apache.org/>


 PlaidML

❑ **PlaidML** - open source portable deep learning engine from Vertex.AI.
<http://vertex.ai/blog/announcing-plaidml>
<https://github.com/plaidml/plaidml>

❑ **A.I. Model** - a common machine learning tool for all frameworks by Microsoft.
<https://www.microsoft.com/en-us/AI/ai-platform>

three major tools for developers:

- Azure Machine Learning Experimentation Service
- Azure Machine Learning Workbench
- Azure Machine Learning Model Management Service.




❑ **IBM Watson** and corresponding Watson Data Platform and IBM Services Platform.
<https://www.ibm.com/watson/>



IBM Watson

❑ **Google Cloud AI** provides modern machine learning services, with pre-trained models and a service to generate your own tailored models.
<https://cloud.google.com/products/machine-learning/>

❑ **Intel® AI DevCloud** for Intel® AI Academy members.
<https://software.intel.com/ai-academy/tools/devcloud>



Cognitive Computing Services



Google (DeepMind)

<https://cloud.google.com/vision/>
<https://cloud.google.com/speech/>

...



IBM Watson

IBM Watson

<https://www.ibm.com/watson/services/visual-recognition/>
<https://www.ibm.com/watson/services/text-to-speech/>
<https://www.ibm.com/watson/services/natural-language-understanding/>

...



Microsoft

Microsoft Azure

<https://azure.microsoft.com/en-us/services/cognitive-services/computer-vision/>
<https://azure.microsoft.com/en-gb/services/cognitive-services/speech/>
<https://azure.microsoft.com/en-us/services/cognitive-services/face/>
<https://azure.microsoft.com/en-us/services/cognitive-services/text-analytics/>

...

<https://www.captionbot.ai/>
<https://www.how-old.net/>
<https://www.what-dog.net/>



Amazon

<https://aws.amazon.com/rekognition/>



Meta AI

<https://ai.meta.com/resources/>



Cognitec

<http://www.cognitec.com/>



Clarifai

<https://www.clarifai.com/>



Kairos

<https://www.kairos.com/>



MetaMind

MetaMind

<https://einstein.ai/>
<https://metamind.readme.io/>

Relevant links:

<https://blog.filestack.com/thoughts-and-knowledge/comparing-google-vision-microsoft-cognitive-amazon-rekognition-clarifai/>