

An Introduction to Machine Learning

Heli Helskyaho
@helifromfinland

Introduction, Heli

- * Graduated from University of Helsinki (Master of Science, computer science), currently a doctoral student, researcher and lecturer (databases, Big Data, Multi-model Databases, methods and tools for utilizing semi-structured data for decision making) at University of Helsinki
- * Worked with Oracle products since 1993, worked for IT since 1990
- * Data and Database!
- * CEO for Miracle Finland Oy
- * Oracle ACE Director, Oracle Groundbreaker Ambassador
- * Listed as one of the TOP 100 influences on IT sector in Finland (2015, 2016, 2017, 2018, 2019)
- * Public speaker and an author
- * Author of the book Oracle SQL Developer Data Modeler for Database Design Mastery (Oracle Press, 2015), co-author for Real World SQL and PL/SQL: Advice from the Experts (Oracle Press, 2016)



Oracle SQL Developer Data Modeler for Database Design Mastery

Design, Deploy, and Maintain World-Class Databases
on Any Platform

Heli Helskyaho
Oracle ACE Director

Forewords by C.J. Date and Tom Kyte



Real World SQL & PL/SQL

Advice from the Experts

Arup Nanda
Brendan Tierney
Heli Helskyaho
Martin Widlake
Alex Nuijten



What is Machine Learning?

- * An important part of Artificial Intelligence (AI)
- * Machine learning (ML) teaches *computers* to learn from *experience (algorithms)*
 - * Learn from data and make predictions
 - * Mathematics, statistics,...
- * “field of study that gives computers the ability to learn without being explicitly programmed“
 - Arthur Samuel, 1959
- * A systematic study of algorithms and systems that improve their *knowledge* or *performance* with *experience*

Challenges in Big Data

- * More and more data (volume)
- * Different data models and formats (variety)
- * Loading in progress while data exploration going on (velocity)
- * Not all data is reliable (veracity)
- * We do not know what we are looking for (value, viability, variability)
- * Must support also non-technical users (journalists, investors, politicians,...)
(visualization)
- * All must be done *efficiently and fast and as much as possibly by machines*
- * -> MACHINE LEARNING?

Why ML? Why now?

- * It is now possible to use machine learning
 - * Improved technology (GPU, TPU,...)
 - * The price for storage solutions
 - * Existing data sets
 - * Pre-trained models
 - * Research! (FB)
- * ...

Why ML? Why now?

- * An environment that NEEDS ML and is finally able to really use it

When to use ML?

- * You have **data!**
 - * ML cannot be performed without data
- * Rules and equations are
 - * Complex (image recognition)
 - * Constantly changing (fraud detection)
- * The nature of the data changes and the program must adapt (today's spam is tomorrow's ham) (predicting shopping trends)

Real life use cases for ML

- * Spam filters
- * Log filters (and alarms)
- * Data analytics
- * Image recognition
- * Speech recognition
- * Medical diagnosis
- * Robotics
- * ...


Machine Learning

- * Unsupervised learning
 - * finds hidden patterns or intrinsic structures in input data
 - * Clustering
 - * Anomaly detection (can also be supervised or semi-supervised)
- * Supervised learning
 - * trains a model on known input and output data to predict future outputs
 - * Classification
 - * Regression
- * Reinforcement learning
 - * the machine or software agent(s) learns based on feedback from the environment


Unsupervised Learning

- * Learning from unlabeled input data by finding hidden patterns or intrinsic structures in that data, the target is unknown
- * used typically when you
 - * don't have a specific goal
 - * The target variable is unknown
 - * are not sure what information the data contains
 - * want to reduce the features of your data as a preprocessing for supervised learning, dimensionality reduction

Clustering



Anomaly detection




<https://www.qualitydigest.com/inside/metrology-article/anomaly-detection-052919.html>

Supervised Learning


- * Learning from known, labelled data
- * The target is known
- * Training a model on known input and output data to predict future outputs (remember that uncertainty is always involved)
- * Testing the model on known input and output data
 - * to compare the prediction to the known output

Classification


Binary classification:



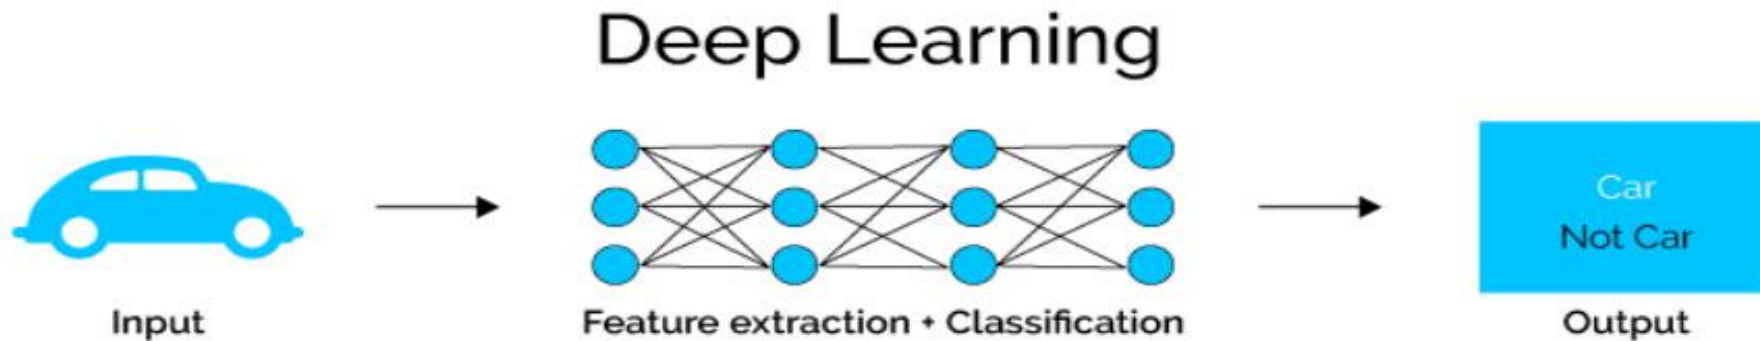

Multi-class classification:



Regression



What is deep learning?

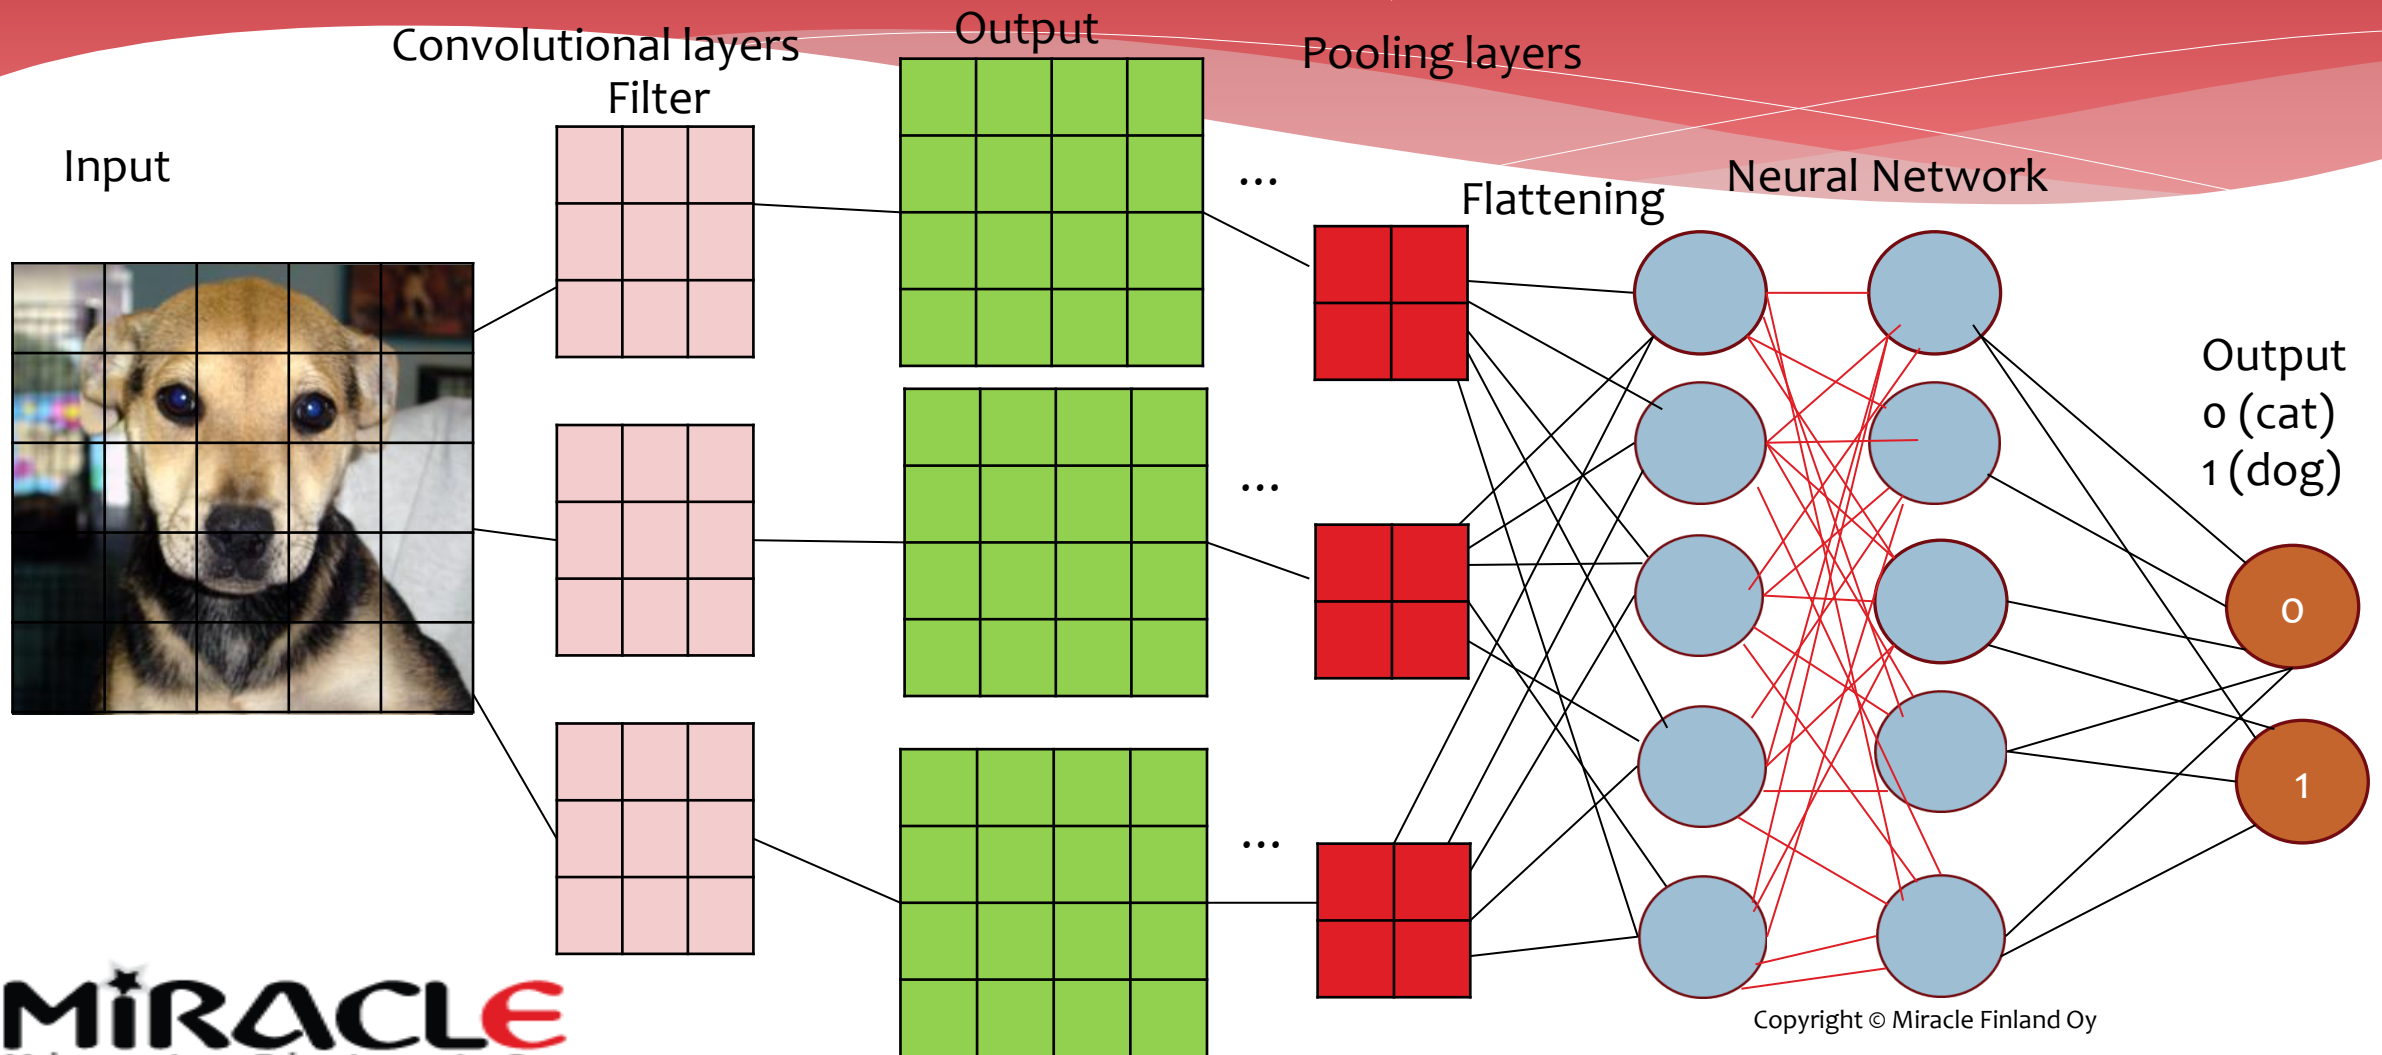


<https://towardsdatascience.com/why-deep-learning-is-needed-over-traditional-machine-learning-1b6a99177063>

Deep Learning (DL)

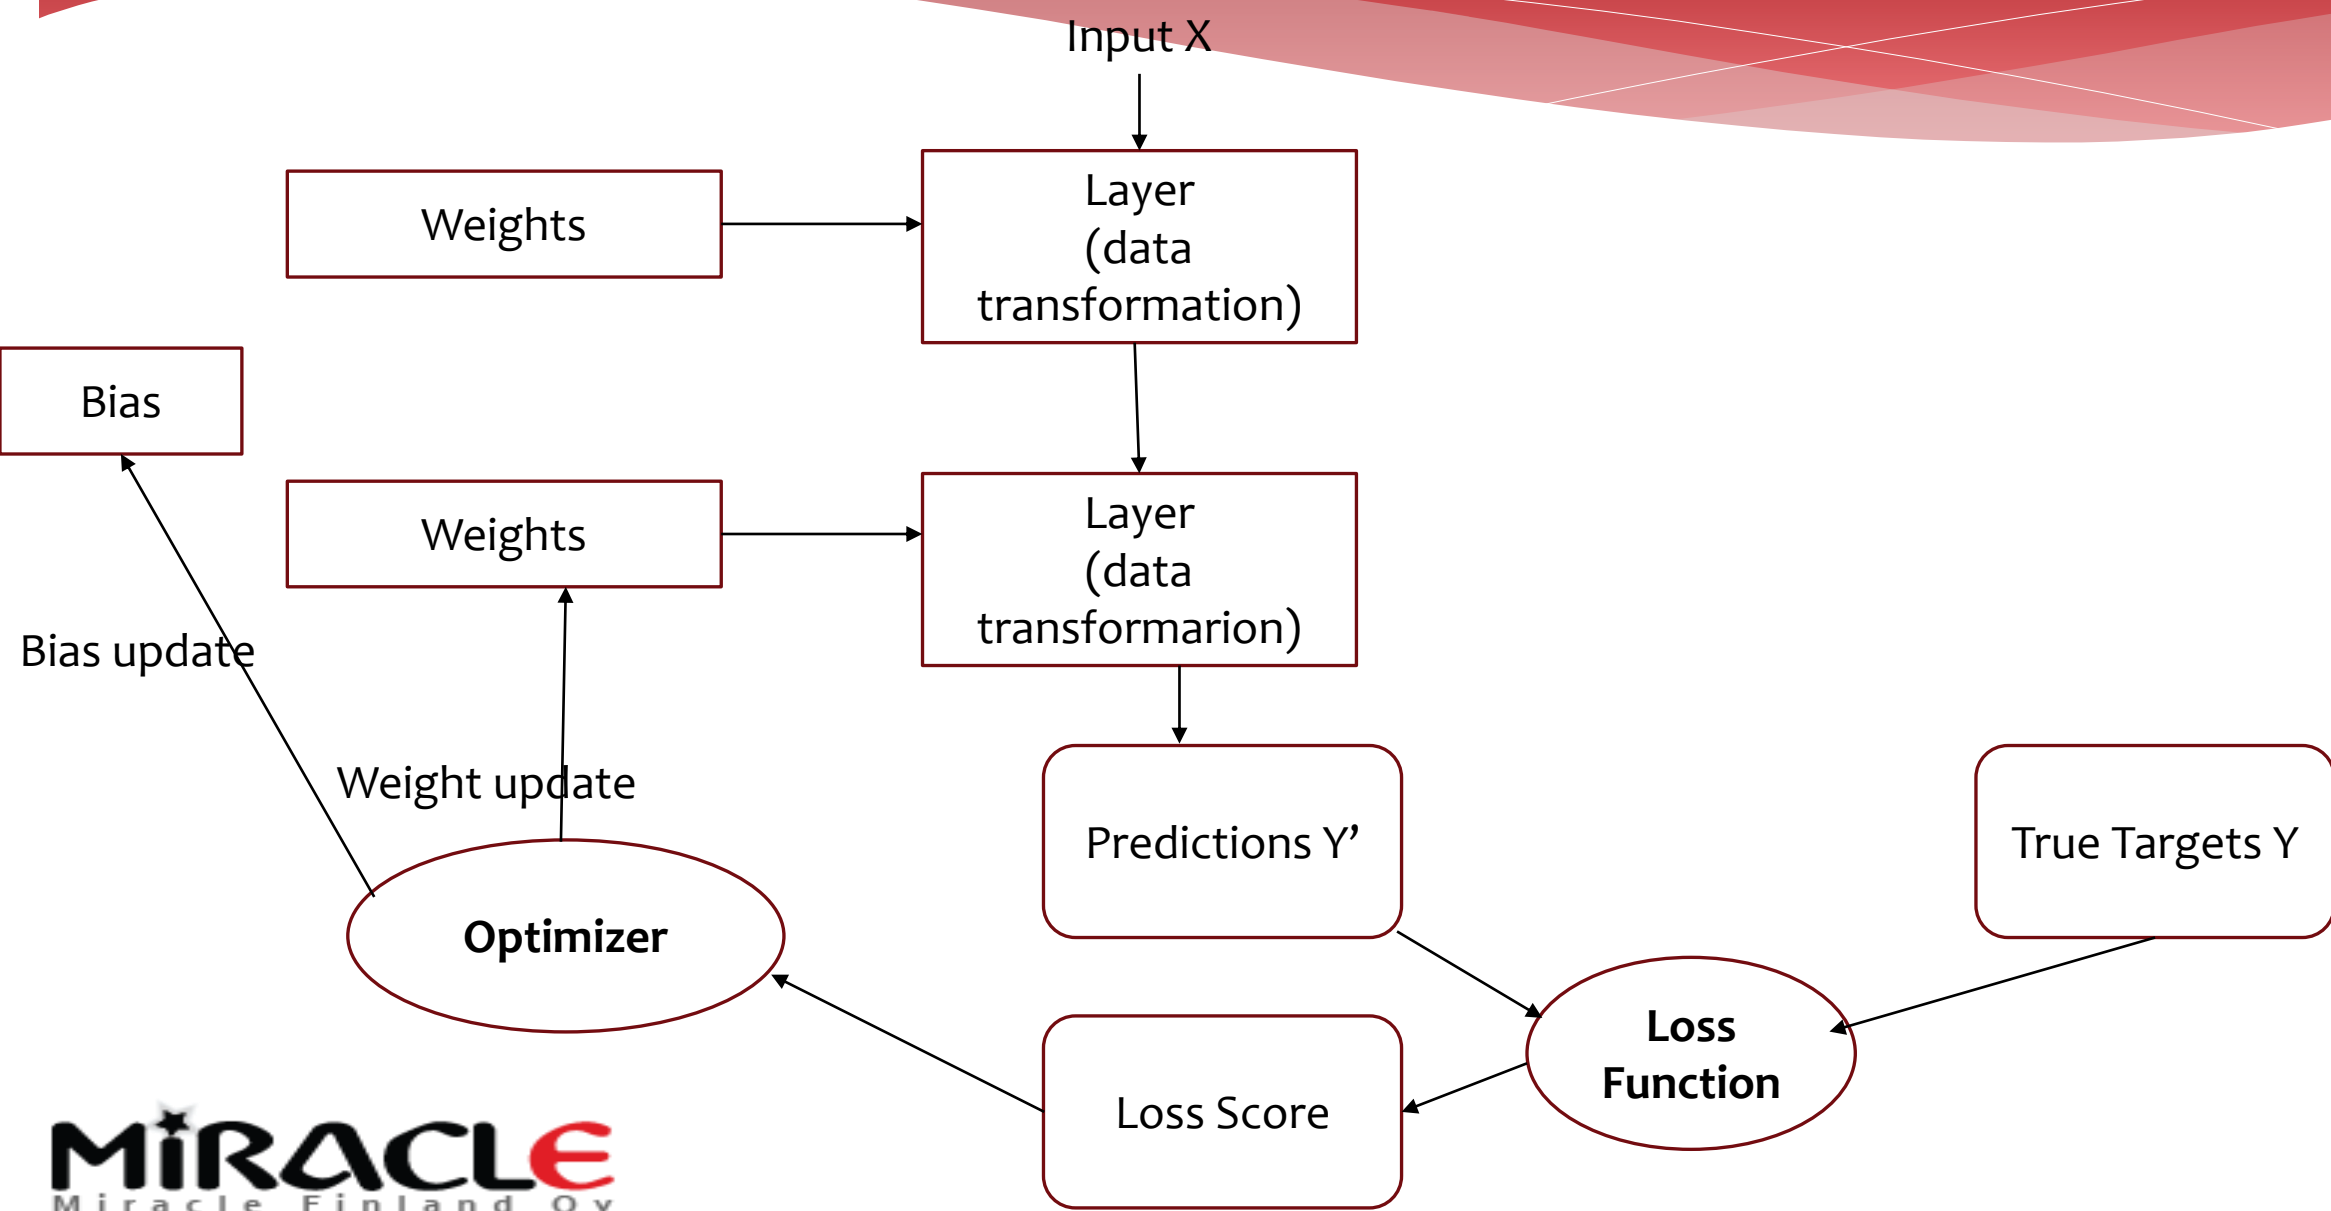
- * Computer vision (Face recognition,...)
- * Speech recognition
- * Text recognition, text processing,...
- * ...
- * Needs a LOT of data and resources
- * Datasets, pre-trained models, transfer learning,...

Convolutional Neural Networks (CNN, or ConvNet)



Backpropagation/backward propagation

- * The process for neural network to learn



Use cases for deep learning?

- * Text, voice, ...
 - * Natural Language Processing
 - * Automatic speech recognition
 - * Text to Speech
 - * Machine Translation
 - * Resume/summary of a text, ...

- * Concerns: "no need to learn to read anymore"

Talk to Transformer

See how a modern neural network completes your text. Type a custom snippet or try one of the examples. [Learn more](#) below.

Custom prompt

Heli Helskyaho will explain machine learning in 25min at Dev Around the Sun event

GENERATE ANOTHER

Completion

Heli Helskyaho will explain machine learning in 25min at Dev Around the Sun event to promote the use of HANA to enable complex machine learning for large scale cloud computing.

Workshops for developers and companies on cloud computing and machine learning technology. Advanced algorithms and machine learning tools. Live tools and demos and information sessions.

New web development career prospects for candidates with dev skillsets. Also mentioned are simpletolearn.com , blog.chai2code.com , tocok.com , ongit.io , irc.freenode.net , and voice.blubbry.com .

There will be posters, tutorials, open projects, courses and a wiki for those who want to make

<https://talktotransformer.com/>

Heli Helskyaho will explain machine learning in 25 min at Dev Around the Sun event at Microsoft Exchange "I want a machine that can make decisions that you don't have time for. It knows how to get an idea, and you can think about a machine like that."

"It's a perfect example because an app cannot talk about this space in just a short time. It needs to be a visual representation if you are talking to a user's desktop.

These AI things are no different. The only way to adapt is to go there. You go like that, I need a visual representation. This is a tool you can't do as a user, I need this image, but I have this image in memory it is the same everywhere. So I can take it and make it do what it is supposed to do. I can't go down it.

But the whole way to learn in AI is to learn how to go down: it's like a tiny piece of paper, it's the little holes that just need to fill. You can think of that same thing as a vector, it takes a really good picture of you and you can choose to use a specific image or tell it how to do things, or that way you can talk to it, or that way you can change.

The best tools that are available if you are building a system like it is the tool. You have to learn how to build a system, and it needs to be adaptable. But the next step is to adapt to that in Artificial Intelligence. In AI you have a bunch of really important things, if

Text Generation API

by Connor Leahy · ❤️ 13 · ➦ share

The text generation API is backed by a large-scale unsupervised language model that can generate paragraphs of text. This transformer-based language model, based on the GPT-2 model by OpenAI, intakes a sentence or partial sentence and predicts subsequent text from that input.

Heli Helskyaho will explain machine learning in 25 min at Dev Around the Sun event

Submit

API Docs



Translate

Google Käytäjä



Kirjautu sisään

Teksti

Dokumentit

TUNNISTA KIELI

SUOMI

ENGLANTI

RUOTSI



SUOMI

ENGLANTI

RUOTSI



Heli pitää esityksen syväoppimisesta APEX
Connect-tapahtumassa.



Heli will give a presentation on in-depth learning at
APEX Connect.



63/5000







Use cases for deep learning?

- * Recommender systems
 - * produce suggestions/recommendations to assist their users in decision-making processes
 - * Collaborative filtering (CF) recommender systems: recommendations based on decisions of *other users* with similar tastes
 - * Content-based recommender systems: recommendations based on similarities of new *items* to those that the user liked in the past
 - * Hybrid recommender systems: *multiple* approaches together

Amazon

Kindle book deals for you



[Back to top](#)

[Get to Know Us](#)

[Careers](#)

[Make Money with Us](#)

[Sell on Amazon](#)

[Amazon Payment Products](#)

[Amazon Business Card](#)

[Let Us Help You](#)

[Amazon and COVID-19](#)

Netflix

NETFLIX

Etusivu

Sarjat

Elokuvat

Uusimmat

Oma lista



Koska katsoit tämän: Another Life



Ajanvietettä




Suosikkeja sinulle, Heli




Use cases for deep learning?

- * Visual recognition problems, Computer vision:
 - * Image classification
 - * Object recognition
 - * Object detection
 - * Image captioning
 - * Action classification
 - * Object/image segmentation (separate the object from its background)
 - * ...





What if data is not available?

- * The computer creates its own data
- * https://www.youtube.com/watch?v=VMp6pq6_QjI











Throttle: 0,02
Brake: 0
Steering: 1


Attempts: 30004





* 310 000 attempts later: starting to be confident

Reinforcement learning (RL)



<https://towardsdatascience.com/introduction-to-various-reinforcement-learning-algorithms-i-q-learning-sarsa-dqn-ddpg-72a5e0cb6287>

Reinforcement learning (RL)


- * Very good with games
 - * AlphaGo, AlphaGo Zero
 - * <https://deepmind.com/blog/article/alphago-zero-starting-scratch>
- * BUT is it good with everything else?
 - * Needs a lots of try and error... learns slower than humans...
 - * How about teaching the car to park on a real parking lot? Maybe not...

What do we actually want with AI?

- * A machine that can learn (and think) like a human
 - * Deep learning with reasoning (common sense)
 - * Ability to learn background information without attaching it to any particular task (curious, observations)

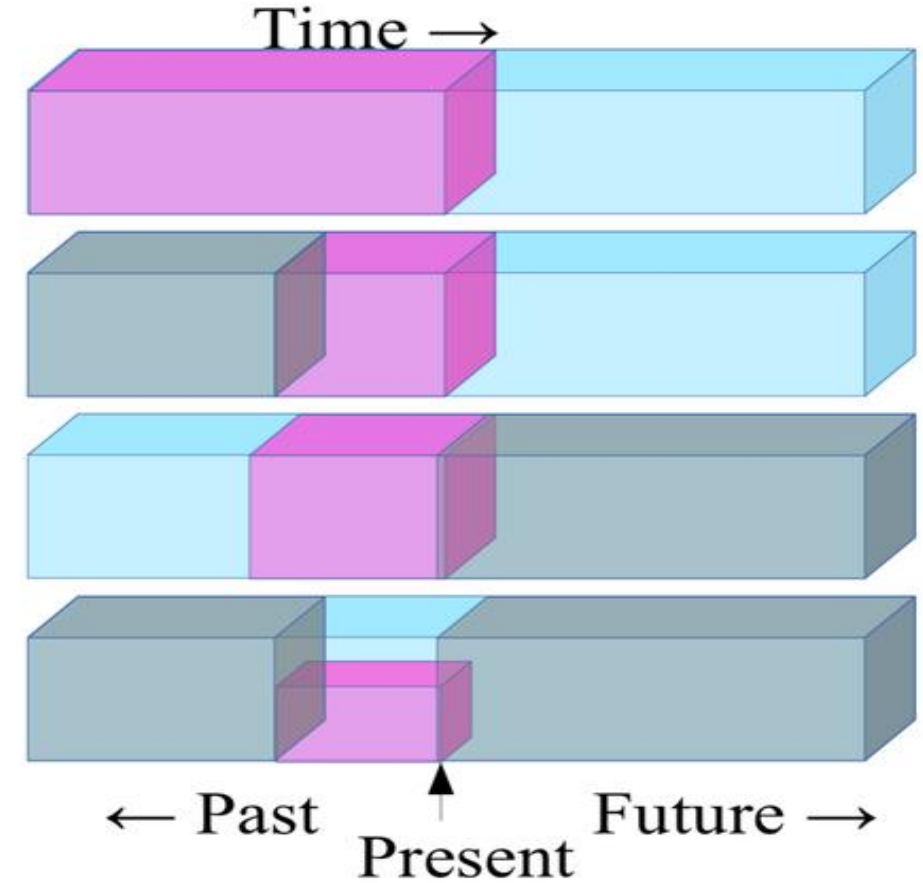
Predictions

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



Self-Supervised Learning

- ▶ Predict any part of the input from any other part.
- ▶ Predict the **future** from the **past**.
- ▶ Predict the **future** from the **recent past**.
- ▶ Predict the **past** from the **present**.
- ▶ Predict the **top** from the **bottom**.
- ▶ Predict the **occluded** from the **visible**
- ▶ **Pretend there is a part of the input you don't know and predict that.**



Self-supervised Reinforcement Learning?

The future and now!

- * AI and machine learning is here and it's the future
- * So many interesting areas to learn
- * Pick your area and **START LEARNING!**

THANK YOU!

QUESTIONS?

Email: heli@miracleoy.fi

Twitter: [@HeliFromFinland](https://twitter.com/HeliFromFinland)

Blog: Helifromfinland.com