Machine learning – why such a big deal, and what is the role of data?

Erik Perjons

DSV, Stockholm University

Stockholms universitet

Why is ML such a big deal?

Why is ML such a big deal?

- According to Brynjolfsson and Mcafee (2017), machine learning (ML) is the **most important general-purpose technology of our era**, similar to previous general-purpose technologies such as steam power, electricity and the combustion engine
- They argue that important general-purpose technologies, such as ML, have a **transformational impact on the economy and business, both directly, but also indirectly by enabling complementary innovations and opportunities**

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Why is ML such a big deal?

Stockholms universitet

There are two reasons why ML is such a big deal:

1) **Humans cannot automate a number of tasks since they cannot explain how to carry out these tasks**. In such cases, **traditional software cannot be developed, since it requires the software developer to understand the tasks at hand in order to code the predefined sequence of instructions**. Read Polanyi: Humans know more than they can tell
2) **On the other hand, ML systems are excellent learners, so with ML we can automate these tasks even if human do not know the sequence of instructions. Why? Since ML can learn from data**

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Why is ML such a big deal?

Stockholms universitet

There are two reasons why ML is such a big deal:

1) **Humans cannot automate a number of tasks since they cannot explain how to carry out these tasks**. In such cases, **traditional software cannot be developed, since it requires the software developer to understand the tasks at hand in order to code the predefined sequence of instructions**. Read Polanyi: Humans know more than they can tell
2) **On the other hand, ML systems are excellent learners, so with ML we can automate these tasks even if human do not know the sequence of instructions. Why? Since ML can learn from data**

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Stockholms universitet

What is ML?

What is ML?

- Brynjolfsson and Mcafee (2017), define Machine Learning (ML) as "the **machine's ability to keep improving its performance without humans having to explain exactly how to accomplish all the task** it's given"

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

What is ML?

- **ML comes in two major categories:**
  - **supervised learning systems**
  - **unsupervised learning systems**

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Stockholms universitet

Supervised learning system

Supervised learning system

- Supervised learning system are **systems that are trained using "labelled" training data in order to categorize new, unlabeled data, based on the training**
- More precisely, **a supervised learning system use a large set of training data showing both the input** (for example, pictures of dogs and cats) **and the correct output** (for example, correct labels for dogs and cats)
- **After the system has been trained with the training data**, **the system should be able to do the mapping between input and output itself. That is, given new input** (new pictures of dogs and cats), **the system should be able to predict the correct output** (label of dog or label of cat)

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Supervised learning system

A somewhat simplified example of how a supervised learning system address a given problem:

- **First, a ML algorithm** and **set of training data** need to be selected that are suitable for the problem to be addressed;
  - **The selected ML algorithm could look like this: y= kx + l** (but represented as code in order to be run in a computer). The algorithm in this case is called ”linear regression”
  - **The training data could consist of one input (x) and one output (y) value, in the form of a value pair (x,y)**, for example, (1,3), (0,2), (-1,0), etc.
- … (cont. next slide)

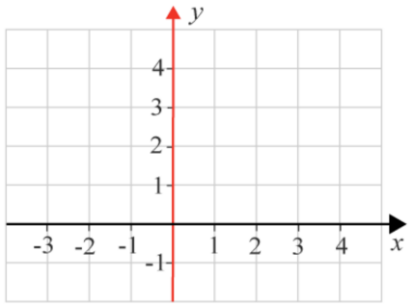

X	Y
1	3
0	2
-1	0

Supervised learning system

- … (from previous slide)
- **Second, the ML algorithm needs to be applied on the training data:**
  - By applying the ML algorithm (y= kx + l) on the training data, the values of the parameters k and l will emerge.
  - For each training data (value pair) used, the algorithm will somewhat adjust the values of the parameters k and l. **This is called training the ML model.**
  - When the ML algorithm has been applied on all training data, **a ML model is created, in this case: y= 2x + 2,** that is, we have received the final values of k and l (k=2 and l=2)
- … (cont. next slide)

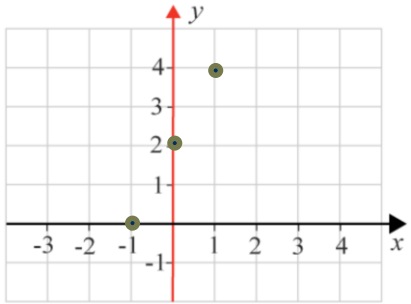

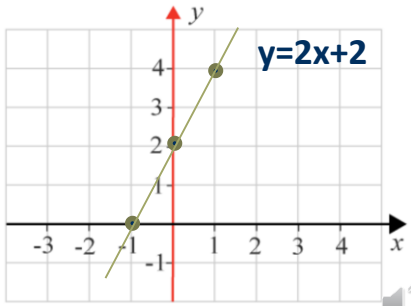

Supervised learning system

- … (from previous slide)
- **Second, the ML algorithm needs to be applied on the training data:**
  - By applying the ML algorithm (y= kx + l) on the training data, the values of the parameters k and l will emerge.
  - For each training data (value pair) used, the algorithm will somewhat adjust the values of the parameters k and l. **This is called training the ML model.**
  - When the ML algorithm has been applied on all training data, **a ML model is created, in this case: y= 2x + 2,** that is, we have received the final values of k and l (k=2 and l=2)
- … (cont. next slide)

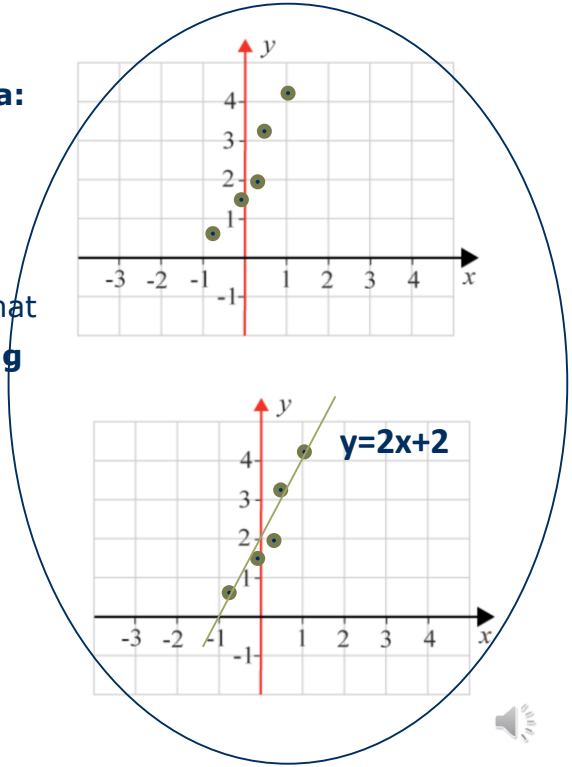

Supervised learning system

- … (from previous slide)
- **Third, this ML model can now be used on new data items as input to predict output:**
  - Each data item has a value for input (x).
  - By applyling the ML model, the value for output (y) will be calculated and predicted

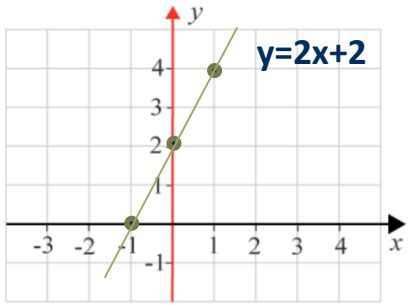

X	Y
3	8
2	6
1	4

Supervised learning system

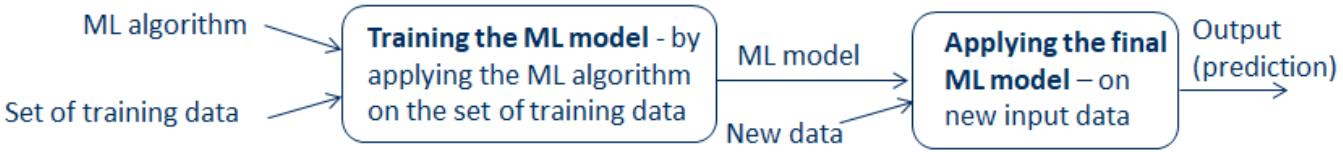

Supervised learning system

- Note, in ML, the use of the terms **”ML algorithm”** and **”ML model”** are sometimes confusing, since they are used in different ways by different authors.
- Other terms used are ”learned function” and ”learning algorithm”

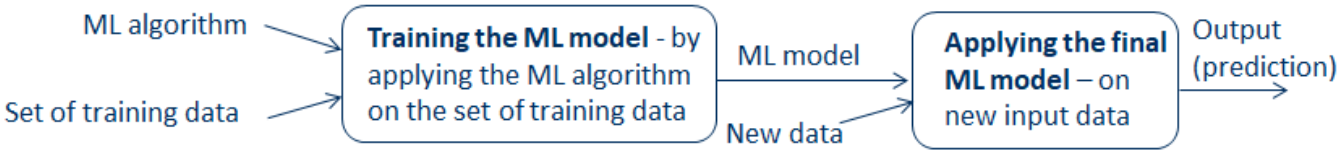

Supervised learning system

- The **selection of an ML algorithm is important** since it **decide the general approach** to take. There are **several guidelines published for selecting an appropriate ML algorithm** to address a problem at hand
- Note also:
  - The **same ML algorithm can be applied** on a **different set of training data**
  - A **different ML algorithm can be applied** on the **same set of training data**

(Alpaydin, E. (2021). Machine Learning, MIT Press Essential Knowledge Series, Chapter 1, s11-33)

Supervised learning systems

- Brynjolfsson & Mcafee (2017) state, referring to Tom Mitchell and Michael I. Jordan, that there are **many examples of recent supervised learning systems** - **which actually map input to output:**

Application	Input X	Output Y
Speech recognition	Voice recording	Transcript
Face recognition	Faces	Name
Customer retension	Purchase histories	Future purchase behavior
Fraud detection	Transaction details	Fraudulent transactions

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Supervised learning systems

- **Supervised learning systems can be applied** when:
  - **the user aims to predict the future** and
  - **the user have a lot of data based on the past (to use for training the system)**
- Examples of applications using supervised learning systems are: diagnosing skin cancer, reviewing loan applications or contracts, and optimizing supply-chains

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Stockholms universitet

Unsupervised learning system

Unsupervised learning systems

- **Unsupervised learning systems aims to learn on their own, without a set of training data.** That is, unsupervised learning systems **finds hidden pattern in unlabeled data**

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Unsupervised learning systems

- **Unsupervised learning systems can discover patterns that we are currently unaware of**
- However, unsupervised learning systems are **difficult to build**
- Note, **humans are excellent unsupervised learners**. Humans are building up their knowledge based on a small amount of data, labeled or not

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Cluster analysis

- **Cluster analysis is a type of unsupervised learning**
- "Cluster analysis or clustering **is the task of grouping a set of objects in such a way that objects in the same group (called a cluster) are more similar (in some sense) to each other than to those in other groups (clusters)**".
- Note, the characteristics of similarity are not known in advance - training data are not used
- An example of cluster analysis is segmenting consumers into similar groups for targeted marketing

(https://en.wikipedia.org/wiki/Cluster_analysis.)

Stockholms universitet

Artificial Neural Network (ANN)

Artificial Neural Network (ANN)

- **A artificial neural network (ANN)** is a machine learning approach that **make use of a network of nodes** that **associate inputs with outputs – via a hidden layer of nodes –** and make use of **weights on these associations** (that is, weights on the associations between the input layer nodes and the hidden layer nodes, and between the hidden layer nodes and the output layer nodes)**, in order to predict output from given input**

- **Each node in an ANN performs some kind of a calculation,** that is, calculation is performed on the receiving input, to be passed on to nodes in another layer

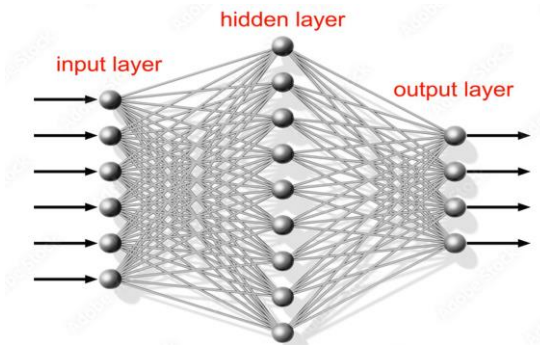

ANN - similar to the brain

- The idea with ANN is to **build algorithms that behave similar to the human brain**
- Nodes in a ANN were inspired by the **neurons** in the human brain. In the brain, these neurons work together in groups connected via synapses. These **synapses** are represented as associations in the ANN

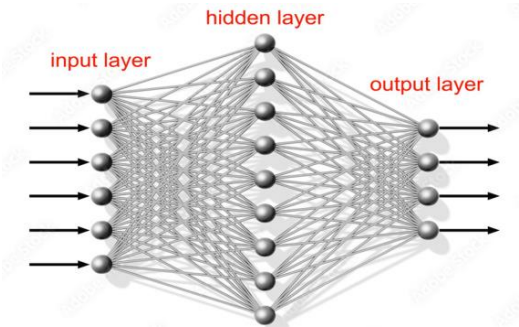

(McCullum, N. (2020) Deep Learning Neural Networks Explained in Plain English, available at (2023-01-9)https://www.freecodecamp.org/news/deep-learning-neural-networks-explained-in-plain-english/)

ANN - weights and activation function

- **Each node in an ANN performs some kind of a calculation** based on the **weight on the association** and an **activation functions**. Then, a new calculated output and sent to nodes in the next layer.
- The data scientist select the weight and the types of activation function to be used

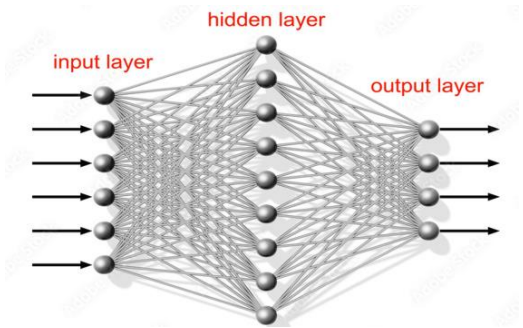

(McCullum, N. (2020) Deep Learning Neural Networks Explained in Plain English, available at (2023-01-9)https://www.freecodecamp.org/news/deep-learning-neural-networks-explained-in-plain-english/)

Training the ANN

- When **training the ANN - the weights are modified** until the **expected output is received** (in this case, it is a form of supervised learning using a set of training data showing the correct output from given input)
- A specific **cost function** is used to train the ANN, that is, to minimize the error of the ANN's prediction. This is done by, as specified above, modifying the weights.

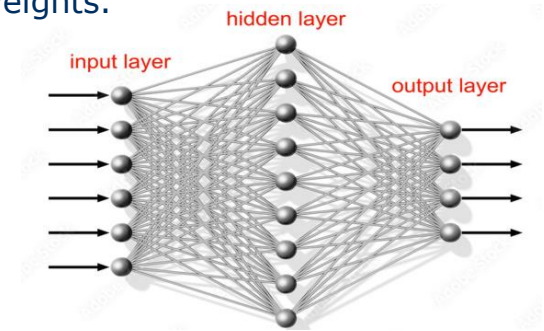

(McCullum, N. (2020) Deep Learning Neural Networks Explained in Plain English, available at (2023-01-9)https://www.freecodecamp.org/news/deep-learning-neural-networks-explained-in-plain-english/)

ANN – a simple example

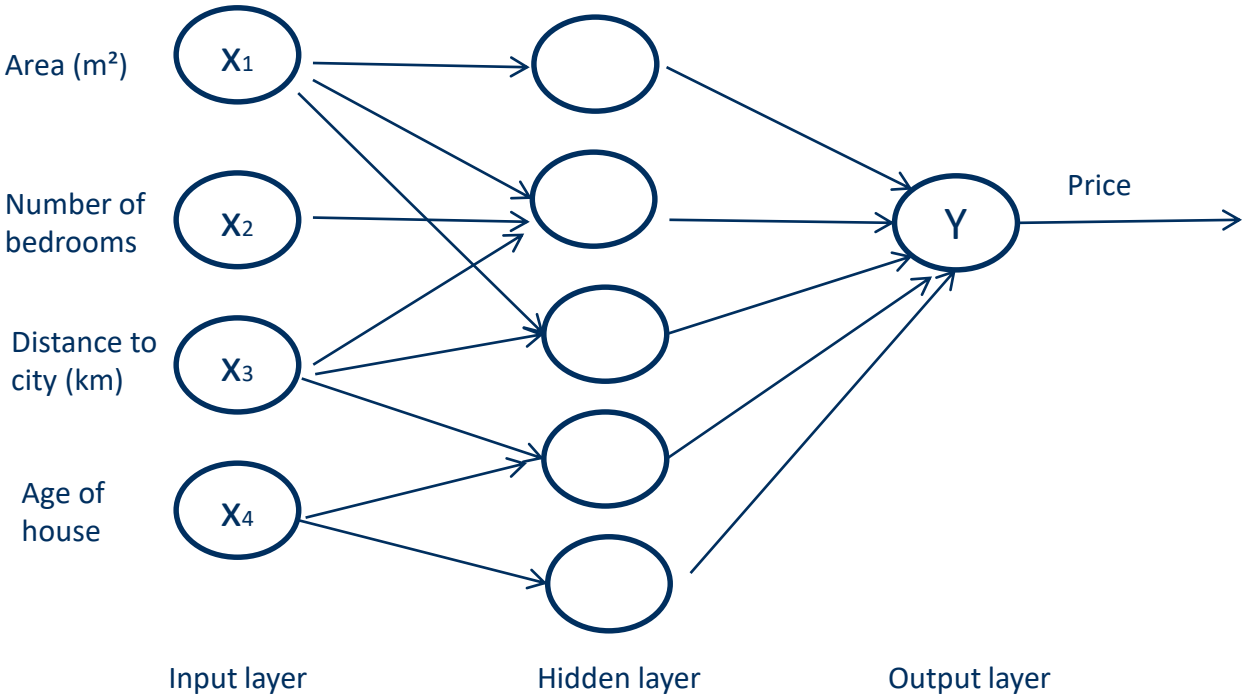

(McCullum, N. (2020) Deep Learning Neural Networks Explained in Plain English, available at (2023-01-9)https://www.freecodecamp.org/news/deep-learning-neural-networks-explained-in-plain-english/)

ANN – a simple example

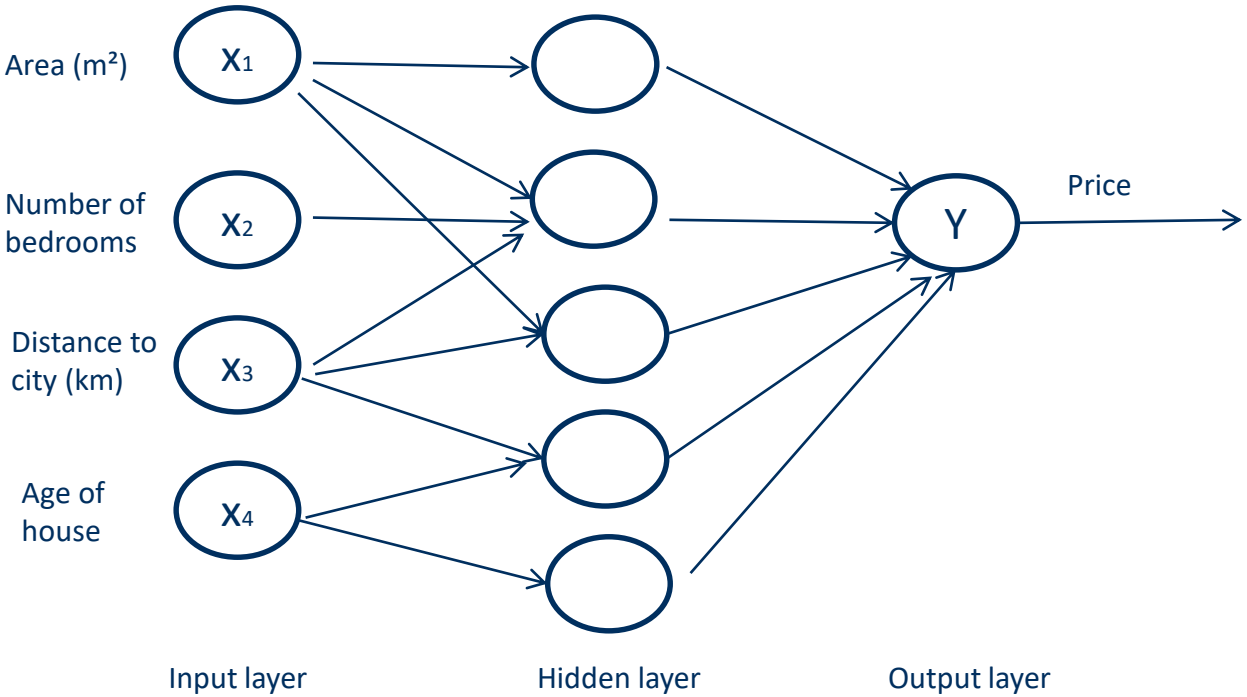

(McCullum, N. (2020) Deep Learning Neural Networks Explained in Plain English, available at (2023-01-9)https://www.freecodecamp.org/news/deep-learning-neural-networks-explained-in-plain-english/)

Stockholms universitet

Deep Learning

Deep learning

- **Deep learning** is a machine learning approach that is using a **"multiple layers neural network"** – that is, an ANN with several hidden layers of nodes
- That is, deep learning is using a so called "deep neural network"

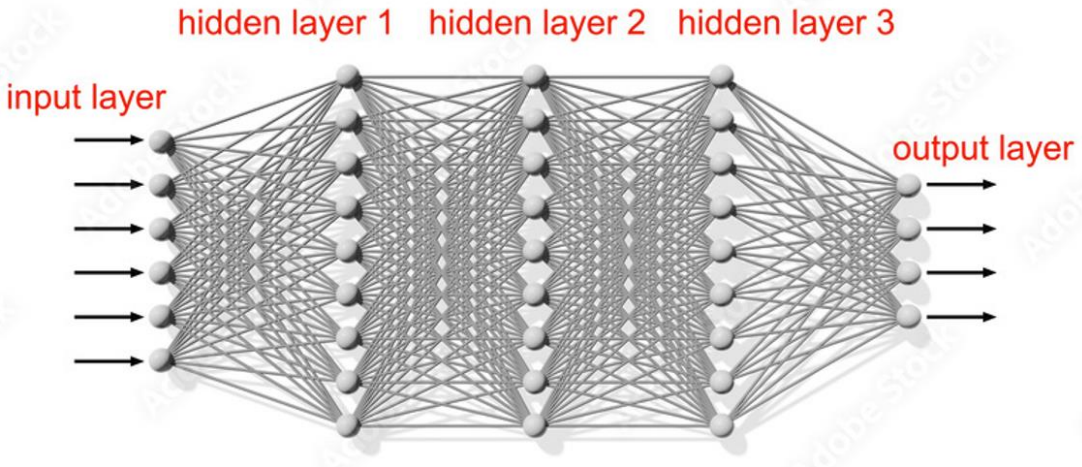

Deep learning

- The **multiple layers of nodes in deep learning** are often **used to progressively extract higher-level features from the raw input**
- "For example, **in image processing, lower layers may identify edges, while higher layers may identify the concepts relevant to a human such as digits or letters or faces**"

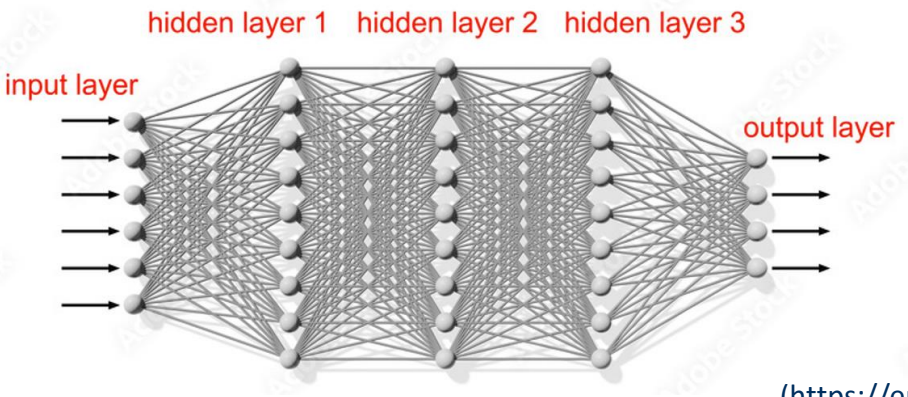

(https://en.wikipedia.org/wiki/Deep_learning)

Deep learning

- **Deep learning has been very successful,** expecially in the areas of **automatic speech** and **image recognition**
- The success of deep learning has more or less **transformed the whole AI industry**

  Therefore, the term ”**deep learning revolution”** is often used when deep learning is discussed
- **Example of deep learning applications:** image recognition, recommender system, financial fraud detection

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Deep learning – benefits

- **Deep learning can make use of larger dataset than traditional ML can do**.
- **More data leads to better predictions using deep learning**, while for traditional ML, more data make better prediction only up to a certain amount of data

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Deep learning - drawbacks

- **Deep neural networks may use hundreds of millions of associations between the nodes**
- **Each of this association contribute to the decision suggested – but each association's impact on the decision is limited**.
- Therefore, deep neural networks **have problem to explain, in a simple way, why it ended up in a certain decision**.
- That is, **deep neural networks knows more than they can tell us** (Note, Polanyi's paradox again)

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Stockholms universitet

Putting ML to work

Putting ML to work

Stockholms universitet

According to Brynjolfsson & Mcafee (2017), **there are three good news for organizations that want to make use of ML**:

- **First, AI skills are spreading quickly**. More and more data scientists are trained and educated
- **Second, algorithms and hardware for ML can be bought or rented, and powerful ML infrastructures are available via the cloud,** all provided by a number of organizations
- **Third**, **to start making productive use of ML, you may not need that much data after all,** if the user accept that significantly improving performance is good enough (and not require to be a dominant actor)

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

ML is driving changes at three levels

According to Brynjolfsson & Mcafee (2017), ML is driving changes at three levels:

- **Task and occupation:** Example of a new task could be to identify potential cancer cells by analyzing radiology images based on ML, carried out by a skilled data scientist
- **Business processes:** Business processes need to change in order to making use of ML, for example, a clinical process within the radiology department need to be changed when using ML for analyzing radiology images.
- **Business models:** Business models may need to be reconsidered when introducing ML solutions, in order to take advantage of these solutions

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Stockholms universitet

Will ML replace jobs?

ML is not replacing the entire job

- According to Brynjolfsson & Mcafee (2017), **most often ML systems do not replace the entire task, process, or business model**
- Instead, **ML systems often complement human activities**, thereby, **making the human's work more efficient and effective**
- **Designing solutions combining ML solutions with humans skills can be very rewarding**, but require human understanding and planning

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Stockholms universitet

Risk and limitation with ML

Three risks with ML

- According to Brynjolfsson & Mcafee (2017), there are three major risks with ML:
  - First, **the ML system may have hidden biases**, derived from the data used for training the system
  - Second, it is **difficult to prove with complete certainty that the system will work in all cases when applying an ML model,** since the ML systems are based on training data and, especially for neural network systems, statistical truth
  - Third, **when the ML system does make errors**, it is difficult to **diagnose and correct exactly what was wrong,** due to the complex structure behind the solution (especially regarding neural network systems)

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

ML limitations

Stockholms universitet

According to Brynjolfsson & Mcafee (2017), ML has its limitations compared to humans:

- ML systems **do not provide general intelligence across diverse domains or different contexts.** That is, ML systems are **trained to do specific tasks, but can typically not be generalized** and do related tasks or similar tasks in another domain
- **Compare ML to human beings**: if a human preforms well in one task, she/he often have competence in related tasks. This not the case for ML
- Moreover, **ML systems are not good in posing questions**, **decide what problem to work on next,** and **decide what area to explore**

(Brynjolfsson, E., & Mcafee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.)

Predicting behavior using ML

Using data to predict behavior

- **Companies want to predict human behavior**, for example, **predict what products the customers are likely to buy in the future**
- **However, predicting customer behavior is tricky since customers change their behavior over time** based on different factors, for example, the customers change behavior based on weather, trends and locations
- **Still we can predict customer behavior. Why?**

(Alpaydin, E. (2021). Machine Learning, MIT Press Essential Knowledge Series, Chapter 1, s11-33)

Using data to predict behavior

- **Predicting customer behavior is possible because: human behavior is not totally random**
- **Customer behavior follow certain patterns**
- **Data, for example sales transactions data, can help us find these patterns, which then can be used for predicting future behavior**

(Alpaydin, E. (2021). Machine Learning, MIT Press Essential Knowledge Series, Chapter 1, s11-33)

ML – to find patterns in data

- **To find the patterns in the data**, an **ML algorithm can be selected and used** in order to **create an ML model.**
- **The ML model could be** seen as a **general template with parameters modified during the training of the ML model using the ML algorithm**

(Alpaydin, E. (2021). Machine Learning, MIT Press Essential Knowledge Series, Chapter 1, s11-33)

The brain and ML

The brain – similar to computers

- According to Marr (1982), **an information processing system**, like the brain and the computer, can be **analyzed on three different levels**:
  - **Computational theory** – defines goal of the system and the overall task to be carried out
  - **Representation and algorithm** – defines how input and output are represented and specify the algorithm for transforming input to output
  - **Hardware implementation** – defines the physical realization of the system

(Alpaydin, E. (2021). Machine Learning, MIT Press Essential Knowledge Series, Chapter 1, s11-33)

The brain – similar to computers

- **The human brain is – in some respect - similar to a computer – and in other not**
- For example, when **human and computer are carrying out calculations:**
  - **the computational theory is the same,**
  - **the representation and algorithm are different**
  - **the hardware implementation is also different**

(Alpaydin, E. (2021). Machine Learning, MIT Press Essential Knowledge Series, Chapter 1, s11-33)

The brain – can be of help in building software solutions

- **Still, the brain can help humans building better computer system**
- More precisely, **human can learn from how the brain make use of the representation and the algorithm in order to design another representation and algorithm, adapted for computers, which is the case when developing artificial neural networks (ANN)**

(Alpaydin, E. (2021). Machine Learning, MIT Press Essential Knowledge Series, Chapter 1, s11-33)

The brain – differs from computers

- **However, the brain and computer are also different in many ways**:
  - **The brain has a very large number of processing units**, called neurons, while the computer only has a few
  - **Each processing units in a brain is much simpler and slower** than a processing units in a computer
  - **The brain has a large connectivity that provides its processing power**. A neuron have connections, called synapses, to ten of thousands other neurons.
  - **In the brain, processing and memory are distributed together over the network of neurons and synapses**. Processing is done by the neurons, but the memory occurs in the synapses between the neurons. In a computer, the processing units and memory are separated

(Alpaydin, E. (2021). Machine Learning, MIT Press Essential Knowledge Series, Chapter 1, s11-33)