MO434 - Deep Learning
Introduction to Deep Learning

Alexandre Xavier Falcão

Institute of Computing - UNICAMP

afalcao@ic.unicamp.br

Agenda

- Artificial Intelligence, Machine Learning, and Deep Learning.
- Deep learning and other areas.
- Objectives of this course.
- What you will learn.
- How your performance will be evaluated.
- Syllabus of the course.
- Introducing PyTorch.

Artificial Intelligence

Artificial Intelligence (1950s) is the part of Computer Science whose goal is to enable machines to behave like humans in

- reasoning,
- representing knowledge,
- planing,
- learning,
- interpreting a natural language,
- moving, perceiving, understanding, manipulating, interacting, etc.

Machine Learning

Machine Learning is the subfield of AI whose goal is to learn from data (text, image, voice) – digital information acquired from the environment with/without extra knowledge transfered by a human.

Machine Learning

Machine Learning is the subfield of AI whose goal is to learn from data (text, image, voice) – digital information acquired from the environment with/without extra knowledge transfered by a human.

- In classical computer programming, rules and data produce answers.

  For instance, given the data, mass x_1 of an object and acceleration x_2 of its motion, the rule $y = x_1 x_2$ finds the force y (answer) that causes the motion.

Machine Learning

Machine Learning is the subfield of AI whose goal is to learn from data (text, image, voice) – digital information acquired from the environment with/without extra knowledge transfered by a human.

- In classical computer programming, rules and data produce answers.

  For instance, given the data, mass x_1 of an object and acceleration x_2 of its motion, the rule $y = x_1 x_2$ finds the force y (answer) that causes the motion.

- In machine learning, data and answers produce rules – i.e., a model to predict answers from new data.

  For instance, y may be the price (answer) of a house in a city and the data $\mathbf{x} = (x_i)_{i=1}^{n}$ (a vector) may contain distances to recreation areas, size, number of bedrooms, bathrooms, and appliances, actual and apparent ages, etc. The rule F is a computational model to perform $y = F(\mathbf{x})$.

- The model F may predict measures $\mathbf{y} = (y_j)_{j=1}^{k}$ from observations $\mathbf{x} = (x_i)_{i=1}^{n}$, being called a classifier, when $y_j \in \{0, 1\}$, and a regressor, when the y_j's are continuous values.

Machine Learning

- The model F may predict measures $\mathbf{y} = (y_j)_{j=1}^{k}$ from observations $\mathbf{x} = (x_i)_{i=1}^{n}$, being called a classifier, when $y_j \in \{0, 1\}$, and a regressor, when the y_j's are continuous values.

- When the answers are fully (partially) provided by a human, the process is said supervised (semi-supervised). Otherwise, it is unsupervised.

Machine Learning

- The model F may predict measures $\mathbf{y} = (y_j)_{j=1}^{k}$ from observations $\mathbf{x} = (x_i)_{i=1}^{n}$, being called a classifier, when $y_j \in \{0, 1\}$, and a regressor, when the y_j's are continuous values.

- When the answers are fully (partially) provided by a human, the process is said supervised (semi-supervised). Otherwise, it is unsupervised.

- A fourth paradigm is reinforcement learning, in which data and answers are acquired from the experience of the intelligent agent (apprentice model) in the environment. Human input and methods from the three other paradigms can be used.

Exercise

Other paradigms are:

- meta, tranfer, few-shot, and zero-shot learning,
- continuous (online), multi-task, and active learning,
- contrastive and self learning,
- multiple instance and weakly supervised learning.

Some of them will be used in this course.

Deep Learning

Deep Learning is the ML methodology whose goal is to learn data representation by a sequence of layers (data transformations) – higher is the number of layers, deeper is the model.

Deep Learning

Deep Learning is the ML methodology whose goal is to learn data representation by a sequence of layers (data transformations) – higher is the number of layers, deeper is the model.

- The data is a set X with N examples (samples) of $\mathbf{x} = (x_i)_{i=1}^{n}$.

Deep Learning

Deep Learning is the ML methodology whose goal is to learn data representation by a sequence of layers (data transformations) – higher is the number of layers, deeper is the model.

- The data is a set X with N examples (samples) of $\mathbf{x} = (x_i)_{i=1}^{n}$.
- For each sample, its answer (true target) $\mathbf{y} = (y_j)_{j=1}^{k}$ may also be provided for supervised learning.

Deep Learning

Deep Learning is the ML methodology whose goal is to learn data representation by a sequence of layers (data transformations) – higher is the number of layers, deeper is the model.

- The data is a set X with N examples (samples) of $\mathbf{x} = (x_i)_{i=1}^{n}$.
- For each sample, its answer (true target) $\mathbf{y} = (y_j)_{j=1}^{k}$ may also be provided for supervised learning.
- The model is called a neural network, but the layers may contain other data processing operations.

Deep Learning

Deep Learning is the ML methodology whose goal is to learn data representation by a sequence of layers (data transformations) – higher is the number of layers, deeper is the model.

- The data is a set X with N examples (samples) of $\mathbf{x} = (x_i)_{i=1}^{n}$.
- For each sample, its answer (true target) $\mathbf{y} = (y_j)_{j=1}^{k}$ may also be provided for supervised learning.
- The model is called a neural network, but the layers may contain other data processing operations.
- The first layers, called encoder (backbone), obtain a suitable high-dimensional and sparse data representation.

Deep Learning

Deep Learning is the ML methodology whose goal is to learn data representation by a sequence of layers (data transformations) – higher is the number of layers, deeper is the model.

- The data is a set X with N examples (samples) of $\mathbf{x} = (x_i)_{i=1}^{n}$.
- For each sample, its answer (true target) $\mathbf{y} = (y_j)_{j=1}^{k}$ may also be provided for supervised learning.
- The model is called a neural network, but the layers may contain other data processing operations.
- The first layers, called encoder (backbone), obtain a suitable high-dimensional and sparse data representation.
- The last layers perform prediction and may be a regressor (decoder) or a classifier.

Deep Learning

The deep learning process requires an optimizer based on a loss function to adjust the neuron weights in the encoder's and predictor's layers.

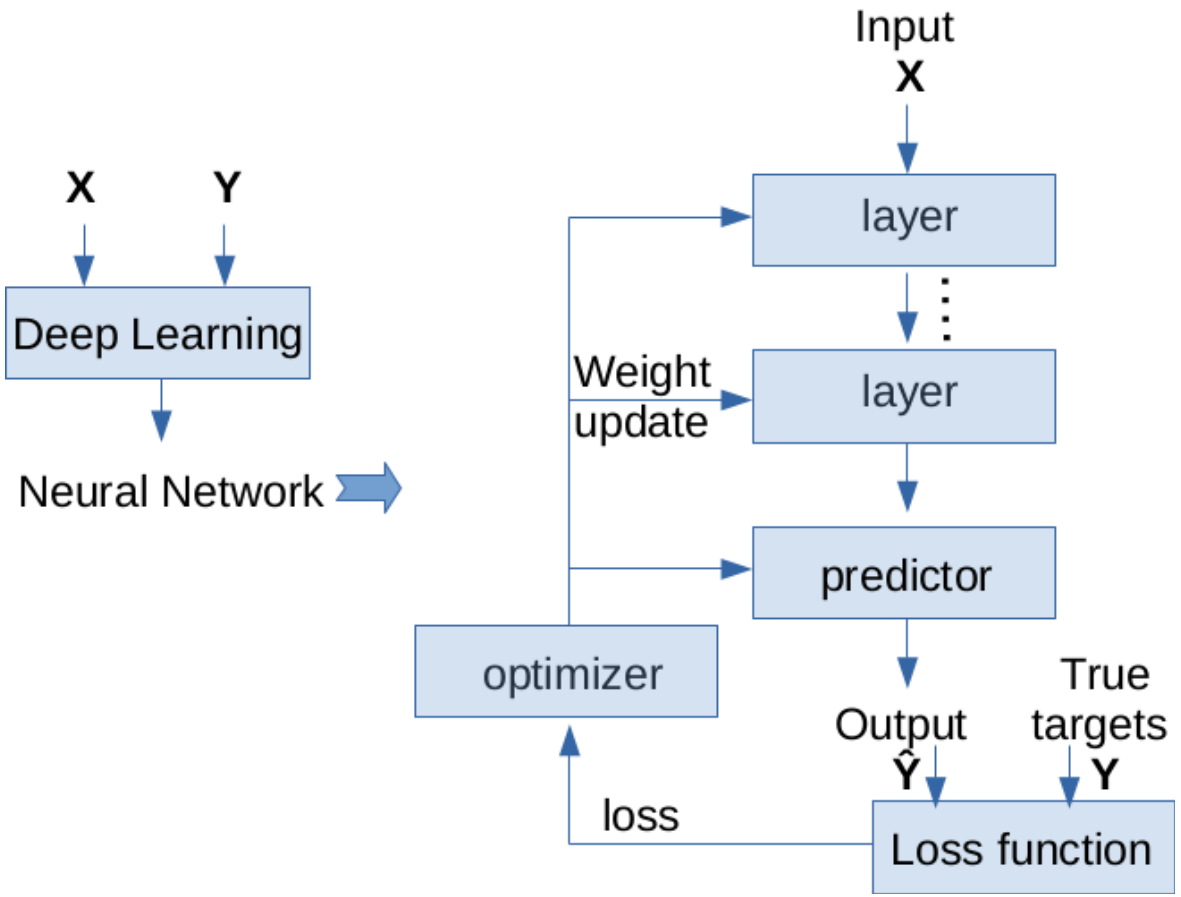

Deep Learning and other areas

DL has applications to Business, Sciences and Engineering. Examples are

- Autonomous cars.
- Virtual assistants.
- Fraud detection.
- Face and speech recognition.
- Medical diagnosis.
- Investment modeling.
- Handwriting transcription.

Objectives

- This course will focus on the design and use of deep neural networks in image and text analysis.
- Prior knowledge on statistics, linear algebra, optimization, machine learning, image and text analysis is important, but the concepts will be introduced whenever they are required.
- The syllabus of the course can be found at `www.ic.unicamp.br/~afalcao/mo434`.

What you will learn

- Deep neural networks (DNNs) for classification and regression.
- The art of training DNNs.
- Fundamentals for image analysis using DNNs.
- Convolutional Neural Networks (CNNs) with Pytorch.
- Classical architectures and information visualization.
- Image classification and related applications with CNNs.
- Fully CNNs (FCNs) for image segmentation.
- Fundamentals for text analysis (nltk, spacy, scrapy).
- Recurrent neural networks and transformers (hugging face).
- Image captioning and other applications in text analysis.

The lectures will be complemented with hands-on activities using jupyter notebooks.

How your performance will be evaluated

You will be graded based on

- your participation during the lectures,
- the practical exercises given at the end of several lecture notebooks, and
- a practical project that will be presented during the course.

Implementation clarity (with comments) and correctness will be considered. A report about the project and exercises shall be presented at the end of the course. For further details, see `www.ic.unicamp.br/~afalcao/mo434`.

Introducing Pytorch

- Pytorch (similar to Tensorflow) provides more flexibility and efficiency than Keras to manipulate data (objects) when building, training and deploying deep neural networks.
- The objects (e.g., images, filters) are represented by Tensors – multidimensional arrays such as ndarrays in numpy.

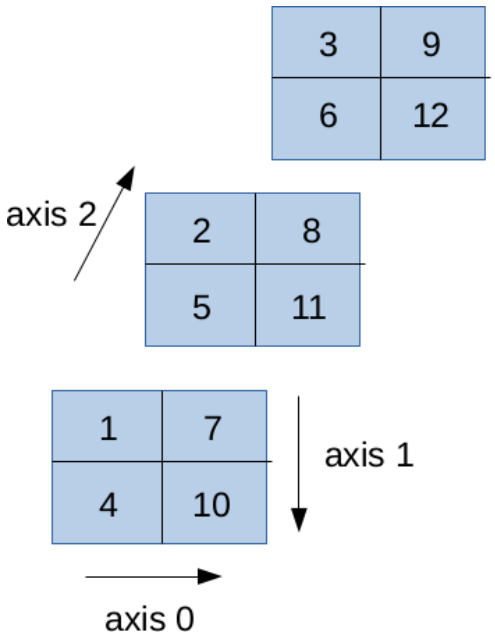

A tensor T1 converted from Numpy ndarray will have T1.shape (2,2,3)

```
tensor([[[ 1,  2,  3],
         [ 4,  5,  6]],

        [[ 7,  8,  9],
         [10, 11, 12]]])
```

T1[1,0,2] is tensor(9)

Introducing Pytorch

- Pytorch (similar to Tensorflow) provides more flexibility and efficiency than Keras to manipulate data (objects) when building, training and deploying deep neural networks.

- The objects (e.g., images, filters) are represented by Tensors – multidimensional arrays such as ndarrays in numpy.

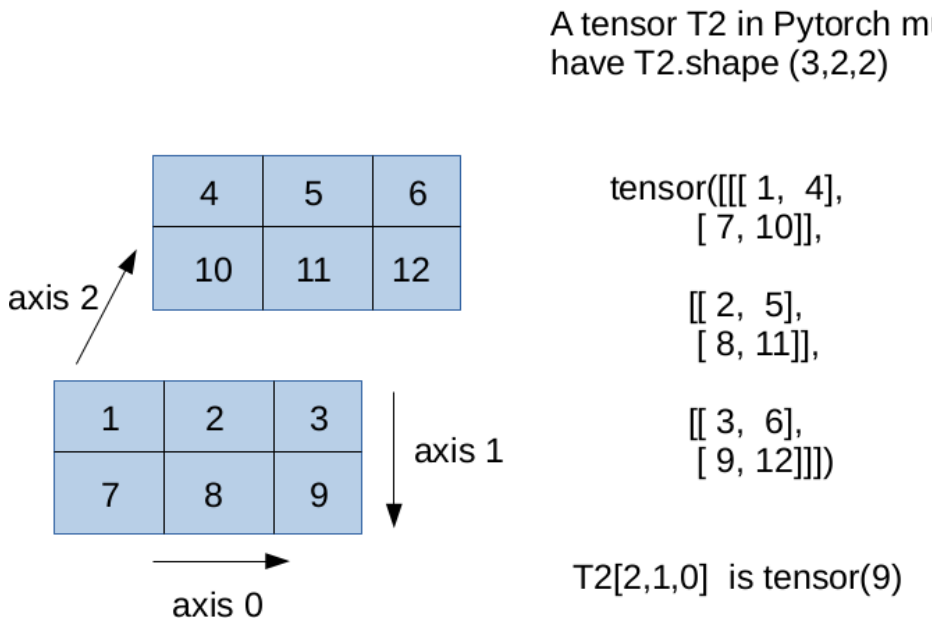

- How to define tensors, manipulate them and implement a trivial NN in Pytorch ▸ Introducing PyTorch.