

DEEP LEARNING 2

FROM THEORY TO PRACTICE

**Matteo Sammut,
PhD Student, Université Paris-Dauphine - PSL**

**Constant Bourdrez,
PhD Student, École Normale Supérieure Paris - PSL**

Double Licence Intelligence Artificielle et Sciences des Organisations
3e année de Licence
Université Paris-Dauphine - PSL

October 7, 2026

Slides originally created by Alexandre Vérine, now at Microsoft AI

SEMESTER SCHEDULE (TEMPORARY)

- ▶ 09/09: Fundamentals of Deep Learning + In a Deep Learning Model
- ▶ 16/09: TP1 Classification - *Introduction to PyTorch*
- ▶ 23/09: In a Deep Learning Model + Techniques to Improve Deep Learning Training + Advanced Deep Learning Techniques
- ▶ 30/09: TP2 Autoencoders - *Hyperparameter Tuning*
- ▶ 07/10: TP3 Image Segmentation - *From CPU to GPU and Parallelization*
- ▶ 14/10: **Graded Individual Practical Work** (To be moved potentially)

SEMESTER SCHEDULE (TEMPORARY)

- ▶ 21/10: TP4 Deep Reinforcement Learning - *From Notebook to Script* - Part 1
- ▶ 28/10: **No Class**
- ▶ 04/11: TP4 Deep Reinforcement Learning - *From Notebook to Script* - Part 2 (To be moved potentially)
- ▶ 11/11: **No Class - Armistice Day**
- ▶ 18/11: TP5 Adversarial Attacks - *Importance of Git* - Part 1
- ▶ 25/11: TP5 Adversarial Attacks - *Importance of Git* - Part 2
- ▶ 02/12: TP6 GANs
- ▶ 09/12: **No Class**
- ▶ 16/12: **Final Exam**

AI 101: FROM FUNDAMENTALS TO DEEP LEARNING

- 1 Introduction to Artificial Intelligence 9**
 - 1.1 Deep Learning in the AI family 9
 - 1.2 Representation Learning 14
- 2 Neural Networks Fundamentals 19**
 - 2.1 Neurons 20
 - 2.2 Layers 22
 - 2.3 Activation Functions 24
- 3 The Multi-layer Perceptron (MLP) 33**
 - 3.1 The first Deep Learning Model 34
 - 3.2 Stochastic Gradient Descent 35
 - 3.3 Back-propagation 38
 - 3.4 Example : Image classification of handwritten digits from A to Z 60

IN A DEEP LEARNING MODEL : FROM NEURAL NETWORKS TO TRANSFORMER MODELS

1	Convolutional Neural Networks	68
1.1	The Two dimensional Convolution	69
1.2	CNN : Convolutional in a network Networks	77
1.3	CNN in practice: CIFAR 10	84
2	Recurrent Neural Networks	104
2.1	Recurrent Block	105
2.2	LSTM and GRU	107
3	Transformer and Attention Mechanism	117
3.1	Self-Attention Mechanism	118
3.2	Transformers Model	123
4	TP2: Build and use an autoencoder	125
4.1	Formal introduction of an autoencoder	125

TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING

- 1 Techniques to Improve Deep Learning Training 132
 - 1.1 Data Augmentation 133
 - 1.2 Learning Rate Scheduling 134
 - 1.3 Early Stopping 135
 - 1.4 Gradient Clipping 136
 - 1.5 Weight Initialization 137
 - 1.6 Regularization 139
 - 1.7 GPU Acceleration 140

DEEP LEARNING AND APPLICATIONS

1	Learning to act with Deep Reinforcement Learning	142
1.1	Deep Q-Learning	142
1.2	The Cheese Game	144
2	Synthetic Data Generation with Generative Adversarial Networks	149
2.1	GANS Models	149
2.2	MNIST Generation	151
3	Sentiment Analysis with Transformers and GRU	154
3.1	Bert	154
3.2	Sentiment Analysis	157
4	Density Estimation with Normalizing Flows	162
4.1	Estimating Density	162
4.2	Normalizing Flows	163
5	Image Segmentation with U-Net	175
5.1	Image Segmentation	175
5.2	U-Net Architecture	178

MANAGING DEEP LEARNING PROJECTS

1	Why Notebooks are not enough	181
2	Modularity	184
3	Local Environment Setup	190
4	Running Experiments	195
5	Monitoring Experiments	201
6	Versioning your Code	206

Part I

AI 101: FROM FUNDAMENTALS TO DEEP LEARNING

INTRODUCTION TO ARTIFICIAL INTELLIGENCE

DEEP LEARNING IN THE AI FAMILY

In general, among all the class of AI algorithms, we make the difference between 3 sub-categories :

- ▶ **Artificial Intelligence** : human designed program and...
- ▶ **Machine Learning** : human designed features with learned mapping such as Support Vector Machine, Kernels methods, Logistic Regression and ...
- ▶ **Deep Learning**: Learned features with learned mapping such as Multilayer Perceptron, Convolutional Networks, ...




Figure. Subsets of Artificial Intelligence

INTRODUCTION TO ARTIFICIAL INTELLIGENCE




DEEP LEARNING IN THE AI FAMILY

In the field of Artificial Intelligence, the fundamental objective is to find a function f that can perform a desired task. This function can either be set by a human or can be learned through training.

For example, in the context of a binary classification task, the goal is to determine $f(x)$ such that $f(x) = 0$ when the label of x is 0 and $f(x) = 1$ when its label is 1. The choice of AI model impacts the expressivity of the function f .

For example, a logistic regression model uses a linear function to make decisions, where $f(x) = \text{sgn}(Ax + b)$. The expressivity of the model can be increased by using more complex functions, such as polynomials or radial basis functions.

Input data



INTRODUCTION TO ARTIFICIAL INTELLIGENCE

CLASSIFICATION TASK




Figure. 2D classification for different AI models.

INTRODUCTION TO ARTIFICIAL INTELLIGENCE

THE UNIVERSAL APPROXIMATION THEOREM

The Universal Approximation Theorem is a fundamental result in the field of artificial neural networks. It states that a deep learning model can approximate any function.

Theorem 1 (Universal Approximation Theorem)

Let $\mathcal{X} \subset \mathbb{R}^d$ be compact, $\mathcal{Y} \subset \mathbb{R}^m$, $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a continuous function and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous real function. Then σ is not polynomial if and only if for every $\epsilon > 0$, there exist $k \in \mathbb{N}$, $A \in \mathbb{R}^{k \times d}$, $b \in \mathbb{R}^k$ and $C \in \mathbb{R}^{m \times k}$ such that

$$\sup_{x \in \mathcal{X}} \|f(x) - g(x)\| \leq \epsilon$$

where $g(x) = C \times \sigma(Ax + b)$.

INTRODUCTION TO ARTIFICIAL INTELLIGENCE

CLASSIFICATION TASK

Multi-layer Perceptron







Figure. 2D classification for small Neural Network.

INTRODUCTION TO ARTIFICIAL INTELLIGENCE

REPRESENTATION LEARNING

How does deep learning work in practice ?

Deep learning is a subset of representation learning that uses deep neural networks to learn meaningful representations of data. In deep learning, representations are learned through a hierarchy of nonlinear transformations, where each layer of the network builds upon the previous one to extract increasingly abstract and higher-level features from the input data.



INTRODUCTION TO ARTIFICIAL INTELLIGENCE

EXAMPLE OF REPRESENTATION LEARNING

Consider the task of recognizing objects in images. A traditional approach would be to hand-engineer features such as edge detectors and color histograms that can be fed into a classifier.

However, with deep learning representation learning, the model learns to automatically discover these features from the data. The network might start by learning simple features such as edges and color blobs in the first layer, then build upon these to learn more complex features such as parts of objects in subsequent layers, until finally, the final layer outputs a probability distribution over classes of objects.

In this way, deep learning of representation enables the model to automatically learn a rich and meaningful representation of the data, without the need for manual feature engineering.




Figure. MNIST

INTRODUCTION TO ARTIFICIAL INTELLIGENCE

EXAMPLE OF REPRESENTATION LEARNING




Figure. MNIST : Layer 0




Figure. MNIST : Layer 1




Figure. MNIST : Layer 2

INTRODUCTION TO ARTIFICIAL INTELLIGENCE

EXAMPLE OF REPRESENTATION LEARNING



Figure. MNIST : Layer 0




Figure. MNIST : Layer 2

INTRODUCTION TO ARTIFICIAL INTELLIGENCE


DEEP LEARNING AND NEURAL NETWORKS

Ok, Deep Learning is a model that learns a good representation of the feature. But how?

- ▶ How does it work ?
- ▶ How can we build a model ?
- ▶ How does it learn ?

NEURAL NETWORKS FUNDAMENTALS


Typically, a neural network is defined as a computational model composed of interconnected nodes, organised into layers, that perform transformations on input data.



NEURAL NETWORKS FUNDAMENTALS

NEURONS


If we consider that the Neural Network is a function $f : \mathbb{R}^d \rightarrow \mathbb{R}^m$:



NEURAL NETWORKS FUNDAMENTALS

NEURONS

For example, with an image dataset, the image can be flattened:



$$\in [0, 1]^{d/2 \times d/2}$$

$$x_0 = [0.00, 0.00, \dots, 0.00, 0.99, 0.07 \dots, 0.00, 0.00] \in [0, 1]^d$$


NEURAL NETWORKS FUNDAMENTALS

LAYERS

A layer i is defined by a matrix $A_i \in \mathbb{R}^{k_{i-1} \times k_i}$, a vector $b_i \in \mathbb{R}^{k_i}$ and a nonlinear function $\sigma_i : \mathbb{R} \mapsto \mathbb{R}$. The transformation made by a layer is:

$$x_i = \sigma_i(A_i x_{i-1} + b_i).$$

The non-linear function σ_i the **activation function**.




NEURAL NETWORKS FUNDAMENTALS

LAYERS

A layer i is defined as a matrix $A_i \in \mathbb{R}^{k_{i-1} \times k_i}$, a vector $b_i \in \mathbb{R}^{k_i}$ and a nonlinear function $\sigma_i : \mathbb{R} \mapsto \mathbb{R}$. The transformation made by a layer is:

$$x_i^k = \sigma_i \left(\sum_{l=1}^{k_i} [A_i]_{l,k} x_{i-1}^l + [b_i]_k \right).$$

The non-linear function σ_i the **activation function**.



NEURAL NETWORKS FUNDAMENTALS

ACTIVATION FUNCTIONS

The activation functions play a crucial role in the implementation of deep neural networks, as they allow them to approximate any continuous function, as stated by the Universal Approximation Theorem. We can list some activation function that are commonly used :

- ▶ Linear
- ▶ Sigmoid
- ▶ Hyperbolic Tangent
- ▶ Rectified Linear Unit (ReLU)
- ▶ Leaky Rectified Linear Unit (Leaky ReLU)
- ▶ Exponential Linear Unit (ELU)
- ▶ Sigmoid-Weighted Linear Unit (Swish)
- ▶ Softmax


NEURAL NETWORKS FUNDAMENTALS

LINEAR

- ▶ Linear activation Function:

$$\sigma(x) = x$$

- ▶ Final activation
- ▶ Use case : Regression




NEURAL NETWORKS FUNDAMENTALS

SIGMOID

- ▶ Sigmoid Function:

$$\sigma(x) = \frac{1}{1 + e^{-x}}$$

- ▶ Final activation
- ▶ Use case : Classification




NEURAL NETWORKS FUNDAMENTALS

SOFTMAX

- Softmax Function:

$$\sigma(x_k) = \frac{e^{x_k}}{\sum_{i=1}^{k_i} e^{x_i}}$$

- Final activation
- Use case : Multi-class Classification




NEURAL NETWORKS FUNDAMENTALS

HYPERBOLIC TANGENT

- Hyperbolic Tangent

$$\sigma(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$$

- Final activation
- Use case : Generative task




NEURAL NETWORKS FUNDAMENTALS

ReLU

- ▶ Rectified Linear Unit (ReLU):

$$\sigma(x) = \max\{0, x\}$$

- ▶ Intermediate activation




NEURAL NETWORKS FUNDAMENTALS

LEAKY ReLU

- ▶ Leaky Rectified Linear Unit (Leaky ReLU):

$$\sigma(x) = \max\{\alpha x, x\}$$

- ▶ Intermediate activation




NEURAL NETWORKS FUNDAMENTALS

ELU

- ▶ Exponential Linear Unit (ELU):

$$\sigma(x) = \begin{cases} \alpha(e^x - 1) & \text{if } x < 0, \\ x & \text{if } x \geq 0. \end{cases}$$

- ▶ Intermediate activation




NEURAL NETWORKS FUNDAMENTALS

SWISH

- Sigmoid-Weighted Linear Unit (Swish):


$$\sigma(x) = \frac{x}{1 + e^{-x}}$$

- Intermediate activation



THE MULTI-LAYER PERCEPTRON (MLP)

Having discussed the structure of a neural network, we will proceed to examine the process of training a model for a specific task. As an illustration, we will consider the example of a Multilayer Perceptron. The two intermediate activation functions are ReLUs and the final activation is a softmax to perform multi-class classification on MNIST. We will consider only 4 classes.




THE MULTI-LAYER PERCEPTRON (MLP)

THE FIRST DEEP LEARNING MODEL

To introduce the training process, we will consider a 3 layers MLP trained to minimise a loss $\mathcal{L}$ over a given a dataset $\mathcal{D}$.

The model f_θ is parameterised by a vector $\theta = \{A_1, A_2, A_3, b_1, b_2, b_3\}$:

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\theta, \mathcal{D})$$



THE MULTI-LAYER PERCEPTRON (MLP)

STOCHASTIC GRADIENT DESCENT

Stochastic gradient descent (SGD) is widely used in deep learning instead of traditional gradient descent due to its efficiency and faster convergence rate. SGD updates the model parameters after computing the gradient of the loss function with respect to each parameter using only a single randomly selected sample. This leads to a faster convergence rate and improved optimization compared to traditional gradient descent, which uses the entire training dataset to compute the gradient at each iteration.

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\theta, \mathcal{D}) = \arg \min_{\theta} \mathbb{E}_{x \sim \mathcal{D}} [l(x, f_{\theta}(x))]$$

THE MULTI-LAYER PERCEPTRON (MLP)

STOCHASTIC GRADIENT DESCENT

Theoretically the algorithm is the following:

Require: Given a loss function l , a dataset $\mathcal{D} = \{x_1, x_2, \dots, x_N\}$ and a learning rate λ

- 1: Initialize parameters θ
- 2: **while** θ has not converged **do**
- 3: **for** $i = 1$ to N **do**
- 4: Randomly select x_i from the dataset
- 5: Compute gradient of the loss with respect to θ : $\nabla_{\theta} l(x_i, f_{\theta}(x_i))$
- 6: Update parameters $\theta = \theta - \lambda \nabla_{\theta} l(x_i, f_{\theta}(x_i))$
- 7: **end for**
- 8: **end while**
- 9: **return** θ

THE MULTI-LAYER PERCEPTRON (MLP)

SGD IN MINI-BATCH

In practice the algorithm is modified to use mini-batches of data instead of single samples. This is done to improve the stability of the optimization process and reduce the variance of the gradient estimates. The algorithm is as follows:

Require: Given a loss function l , a dataset $\mathcal{D} = \{x_1, x_2, \dots, x_N\}$, a learning rate λ and a batch size b

- 1: Initialize parameters θ
- 2: Initialize the number of batches $B = \lfloor \frac{N}{b} \rfloor$
- 3: **while** θ has not converged **do**
- 4: **for** $i = 1$ to B **do**
- 5: Randomly select a mini-batch of b samples from the dataset
- 6: Compute gradient of the loss with respect to θ : $\frac{1}{b} \sum_{i=1}^b \nabla_{\theta} l(x_i, f_{\theta}(x_i))$
- 7: Update parameters $\theta = \theta - \lambda \frac{1}{b} \sum_{i=1}^b \nabla_{\theta} l(x_i, f_{\theta}(x_i))$
- 8: **end for**
- 9: **end while**
- 10: **return** θ


THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

At every step t of the gradient descent, setting a learning rate λ , the parameter θ is updated as:

$$\theta_{t+1} = \theta_t - \lambda \nabla_{\theta} l(f(x_i), y_i)$$

But $\theta = \{A_1, A_2, A_3, b_1, b_2, b_3\}$ and the gradient is computed with respect to each parameter.



THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

First we will consider a single data point x , the loss will depend on the output only: $l(f(x))$.

f is a layered composed function. Let us focus on the last layer:

$$f(x) = x_3 = \sigma_3(A_3x_2 + b_3)$$

Therefore:

$$l(f(x)) = l(\sigma_3(A_3x_2 + b_3))$$

To minimise the loss, we have to act on A_3 , b_3 and x_2 .

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

Let us look at the gradients with respect to A_3 :

$$\begin{aligned}\frac{\partial l}{\partial A_3} &= \frac{\partial l}{\partial x_3} \frac{\partial x_3}{\partial A_3} = l'(x_3) \frac{\partial \sigma_3(A_3 x_2 + b_3)}{\partial A_3} = l'(x_3) \sigma'_3(A_3 x_2 + b_3) \frac{\partial [A_3 x_2 + b_3]}{\partial A_3} \\ &= \underbrace{l'(x_3)}_{\in \mathbb{R}} \underbrace{\sigma'_3(A_3 x_2 + b_3)}_{\in \mathbb{R}^{k_i \times 1}} \underbrace{x_2^T}_{\in \mathbb{R}^{1 \times k_{i-1}}}\end{aligned}$$

and therefore:

$$A_3 \leftarrow A_3 - \lambda l'(x_3) \sigma'_3(A_3 x_2 + b_3) x_2^T.$$

We need to keep in memory the latent values of x , i.e. x_2 .

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

Let us look at the gradients with respect to A_2 :

$$\begin{aligned}\frac{\partial l}{\partial A_2} &= \frac{\partial l}{\partial x_2} \frac{\partial x_2}{\partial A_2} \\ &= \frac{\partial l}{\partial x_2} \frac{\partial \sigma_2 (A_2 x_1 + b_2)}{\partial A_2} \\ &= \frac{\partial l}{\partial x_2} \sigma'_2 (A_2 x_1 + b_2) \frac{\partial [A_2 x_1 + b_2]}{\partial A_2} \\ &= \frac{\partial l}{\partial x_2} \sigma'_2 (A_2 x_1 + b_2) x_1^T\end{aligned}$$

which depends on $\frac{\partial l}{\partial x_2}$, we need to compute it.

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

We have to compute the gradient with respect to x_2 :

$$\begin{aligned}\frac{\partial l}{\partial x_2} &= \frac{\partial l}{\partial x_3} \frac{\partial x_3}{\partial x_2} = l'(x_3) \frac{\partial \sigma_3(A_3 x_2 + b_3)}{\partial x_2} = l'(x_3) \frac{\partial [A_3 x_2 + b_3]}{\partial x_2} \sigma'_3(A_3 x_2 + b_3) \\ &= l'(x_3) A_3^T \sigma'_3(A_3 x_2 + b_3)\end{aligned}$$

Therefore:

$$A_2 \leftarrow A_2 - \lambda \left[l'(x_3) A_3^T \sigma'_3(A_3 x_2 + b_3) \times \sigma'_2(A_2 x_1 + b_2) x_1^T \right]$$

The update of A_2 depends on $l'(x_3)$,

BACK-PROPAGATION

We have to compute the gradient with respect to A_1 :

$$\begin{aligned}\frac{\partial l}{\partial A_1} &= \frac{\partial l}{\partial x_1} \frac{\partial x_1}{\partial A_1} \\ &= \frac{\partial l}{\partial x_1} \frac{\partial \sigma_1(A_1 x_0 + b_1)}{\partial A_1} \\ &= \frac{\partial l}{\partial x_1} \sigma'_1(A_1 x_0 + b_0) x_0^T,\end{aligned}$$

which depends on $\frac{\partial l}{\partial x_1}$, we need to compute it.

BACK-PROPAGATION

Let us compute the gradient with respect to x_1 :

$$\begin{aligned}\frac{\partial l}{\partial x_1} &= \frac{\partial l}{\partial x_2} \frac{\partial x_2}{\partial x_1} = \frac{\partial l}{\partial x_2} \frac{\partial \sigma_2(A_2 x_1 + b_2)}{\partial x_1} = \frac{\partial l}{\partial x_2} \frac{\partial [A_2 x_1 + b_2]}{\partial x_1} \sigma'_2(A_2 x_1 + b_2) \\ &= \frac{\partial l}{\partial x_2} A_2^T \sigma'_2(A_2 x_1 + b_2)\end{aligned}$$

Therefore:

$$A_1 \leftarrow A_1 - \lambda \left[l'(x_3) A_3^T \sigma'_3(A_3 x_2 + b_3) A_2^T \sigma'_2(A_2 x_1 + b_2) \times \sigma'_1(A_1 x_0 + b_1) x_0^T \right]$$

BACK-PROPAGATION

In other words, the update on the weights is:

$$A_3 \leftarrow A_3 - \lambda l'(x_3) \sigma'_3 (A_3 x_2 + b_3) x_2^T$$

$$A_2 \leftarrow A_2 - \lambda \left[l'(x_3) A_3^T \sigma'_3 (A_3 x_2 + b_3) \times \sigma'_2 (A_2 x_1 + b_2) x_1^T \right]$$

$$A_1 \leftarrow A_1 - \lambda \left[l'(x_3) A_3^T \sigma'_3 (A_3 x_2 + b_3) A_2^T \sigma'_2 (A_2 x_1 + b_2) \times \sigma'_1 (A_1 x_0 + b_1) x_0^T \right]$$

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

If we look at the update of the different biases, we can easily compute the different gradient and see the updates. First, let us compute the gradient with respect to b_3 :

$$\begin{aligned}\frac{\partial l}{\partial b_3} &= \frac{\partial l}{\partial x_3} \frac{\partial x_3}{\partial b_3} \\ &= l'(x_3) \frac{\partial \sigma_3(A_3 x_2 + b_3)}{\partial b_3} \\ &= l'(x_3) \sigma'_3(A_3 x_2 + b_3) \frac{\partial [A_3 x_2 + b_3]}{\partial b_3} \\ &= \underbrace{l'(x_3)}_{\in \mathbb{R}} \underbrace{\sigma'_3(A_3 x_2 + b_3)}_{\in \mathbb{R}^{k_i \times 1}}\end{aligned}$$

And thus :

$$b_3 \leftarrow b_3 - \lambda l'(x_3) \sigma'(A_3 x_2 + b_3)$$

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

Let's move on the second layer:

$$\begin{aligned}\frac{\partial l}{\partial b_2} &= \frac{\partial l}{\partial x_2} \frac{\partial x_2}{\partial b_2} \\ &= \frac{\partial l}{\partial x_2} \frac{\partial \sigma_2 (A_2 x_1 + b_2)}{\partial b_2} \\ &= \frac{\partial l}{\partial x_2} \sigma'_2 (A_2 x_1 + b_2)\end{aligned}$$

And thus :

$$b_2 \leftarrow b_2 - \lambda \frac{\partial l}{\partial x_2} \sigma'_2 (A_2 x_1 + b_2)$$

We need to back-propagate the term $\frac{\partial l}{\partial x_2}$ computed for the first layer.

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

For the first layer:

$$\begin{aligned}\frac{\partial l}{\partial b_1} &= \frac{\partial l}{\partial x_1} \frac{\partial x_1}{\partial b_1} \\ &= \frac{\partial l}{\partial x_1} \frac{\partial \sigma_1 (A_1 x_0 + b_1)}{\partial b_1} \\ &= \frac{\partial l}{\partial x_1} \sigma'_1 (A_1 x_0 + b_0)\end{aligned}$$

And thus :

$$b_1 \leftarrow b_1 - \lambda \frac{\partial l}{\partial x_1} \sigma'_1 (A_1 x_0 + b_1)$$

We need to back-propagate the term $\frac{\partial l}{\partial x_1}$ computed for the second layer which has been computed with $\frac{\partial l}{\partial x_2}$ back-propagated from the first layer.

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION



To update the weights, we need to compute the gradient of the loss with respect to the output of the network, and then **back-propagate** the gradient of the loss with respect to each activation, the $\frac{\partial l}{\partial x_i}$, through the network to compute the gradients with respect to the weights and biases of each layer.

THE MULTI-LAYER PERCEPTRON (MLP)

LAST LAYER

We can plot the current state of the network for a given input.

The red lines show positive values for A_i , the blue lines represent negative values for A_i . The level of transparency is proportional to the previous neurons.





$$x_3^1 = \sigma_3 \left(A_3^{1,1} x_2^1 + A_3^{1,2} x_2^2 + \cdots + A_3^{1,8} x_2^8 \right)$$

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

Iteratively, the neural networks improves its performance.





$$x_3^1 = \sigma_3 \left(A_3^{1,1} x_2^1 + A_3^{1,2} x_2^2 + \cdots + A_3^{1,8} x_2^8 \right)$$

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

Iteratively, the neural networks improves its performance.





$$x_3^1 = \sigma_3 \left(A_3^{1,1}x_2^1 + A_3^{1,2}x_2^2 + \cdots + A_3^{1,8}x_2^8 \right)$$

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

Iteratively, the neural networks improves its performance.





$$x_3^1 = \sigma_3 \left(A_3^{1,1} x_2^1 + A_3^{1,2} x_2^2 + \cdots + A_3^{1,8} x_2^8 \right)$$

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

Iteratively, the neural networks improves its performance.





$$x_3^1 = \sigma_3 \left(A_3^{1,1} x_2^1 + A_3^{1,2} x_2^2 + \cdots + A_3^{1,8} x_2^8 \right)$$

THE MULTI-LAYER PERCEPTRON (MLP)

LAST LAYER

We can plot the current state of the network for a given input.

Red lines show positive values of A_i , Blue lines represent negative values of A_i . The level of transparency is proportional to the previous neurons.





$$x_3^2 = \sigma_3 \left(A_3^{2,1} x_2^1 + A_3^{2,2} x_2^2 + \cdots + A_3^{2,8} x_2^8 \right)$$

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

Iteratively, the neural networks improves its performance.





$$x_3^2 = \sigma_3 \left(A_3^{2,1} x_2^1 + A_3^{2,2} x_2^2 + \cdots + A_3^{2,8} x_2^8 \right)$$

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

Iteratively, the neural networks improves its performance.





$$x_3^2 = \sigma_3 \left(A_3^{2,1} x_2^1 + A_3^{2,2} x_2^2 + \cdots + A_3^{2,8} x_2^8 \right)$$

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

Iteratively, the neural networks improves its performance.





$$x_3^2 = \sigma_3 \left(A_3^{2,1}x_2^1 + A_3^{2,2}x_2^2 + \cdots + A_3^{2,8}x_2^8 \right)$$

THE MULTI-LAYER PERCEPTRON (MLP)

BACK-PROPAGATION

Iteratively, the neural networks improves its performance.




$$x_3^2 = \sigma_3 \left(A_3^{2,1} x_2^1 + A_3^{2,2} x_2^2 + \cdots + A_3^{2,8} x_2^8 \right)$$

THE MULTI-LAYER PERCEPTRON (MLP)

EXAMPLE : IMAGE CLASSIFICATION OF HANDWRITTEN DIGITS FROM A TO Z

Having discussed the theory behind Artificial Neural Networks and the training process, we will now proceed to demonstrate a comprehensive end-to-end example of image classification on MNIST.



THE MULTI-LAYER PERCEPTRON (MLP)

EXAMPLE : IMAGE CLASSIFICATION OF HANDWRITTEN DIGITS FROM A TO Z



- ▶ Input shape : $1 \times 28 \times 28$.
- ▶ Number of Classes : 10.
- ▶ Number of training samples (x, y) : 60000.
- ▶ Number of evaluating samples: 10000.
- ▶ Loss : cross-entropy

$$L(\hat{y}, y) = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^K y_{ij} \log(\hat{y}_{ij})$$

where :

- $\hat{y} \in \mathbb{R}^{N \times K}$ is the predicted probability distribution over K classes for N samples,
- $y \in 0, 1^{N \times K}$ is the ground-truth one-hot encoded label matrix,

RECAP ON THE CROSS-ENTROPY LOSS



The cross-entropy loss for one sample is:


$$l(\hat{y}_i, y_i) = - \sum_{j=1}^K y_{ij} \log(\hat{y}_{ij}).$$

THE MULTI-LAYER PERCEPTRON (MLP)

EXAMPLE : IMAGE CLASSIFICATION OF HANDWRITTEN DIGITS FROM A TO Z


We build a 3 layers network.

- ▶ Batch size : 64
- ▶ Learning rate : 0.01
- ▶ Intermediate activation : ReLU
- ▶ Final activation : Softmax
- ▶ Number of epochs : 12
- ▶ Number of trained parameters: 52.6k



THE MULTI-LAYER PERCEPTRON (MLP)


EXAMPLE : IMAGE CLASSIFICATION OF HANDWRITTEN DIGITS FROM A TO Z



THE MULTI-LAYER PERCEPTRON (MLP)

EXAMPLE : IMAGE CLASSIFICATION OF HANDWRITTEN DIGITS FROM A TO Z

With a interpretation tool such as SHAP:



Part II

DEEP LEARNING IN ACTION: FROM NEURAL NETWORKS TO TRANSFORMER MODELS

Now that we have an understanding of the training procedure for Artificial Neural Networks, we shall examine several widely-utilized structures within the literature of Neural Networks, including Convolutional Neural Networks (CNN), Residual Networks (ResNet), Recurrent Neural Networks (RNN), and Transformers.

CONVOLUTIONAL NEURAL NETWORKS


In the field of image processing, the Convolution Operators are widely considered as the most favoured approach. While it has been demonstrated that Dense blocks, or Linear blocks, are capable of accurately classifying images in the case of the MNIST dataset, the need for convolutional transformations arises when addressing wider and more intricate datasets.

CONVOLUTIONAL NEURAL NETWORKS

THE TWO DIMENSIONAL CONVOLUTION

A 2D convolution in a neural network context can be mathematically represented as a sliding window operation where a filter (also called kernel) w of size $k \times k$ is applied to each $k \times k$ sub-matrix of the input matrix x . The operation can be defined as the element-wise multiplication of the filter w and the sub-matrix followed by summing the results, i.e.

$$y_{i,j} = \sum_{m=1}^k \sum_{n=1}^k w_{m,n} \cdot x_{i+m,j+n}$$




CONVOLUTIONAL NEURAL NETWORKS

THE TWO DIMENSIONAL CONVOLUTION

A 2D convolution in a neural network context can be mathematically represented as a sliding window operation where a filter (also called kernel) w of size $k \times k$ is applied to each $k \times k$ submatrix of the input matrix x . The operation can be defined as the element-wise multiplication of the filter w and the submatrix followed by summing the results, i.e.

$$y_{i,j} = \sum_{m=1}^k \sum_{n=1}^k w_{m,n} \cdot x_{i+m,j+n}$$



CONVOLUTIONAL NEURAL NETWORKS

KERNEL SIZE, PADDING AND STRIDE

In every Deep Learning library, the Conv2D block takes three parameters in argument:

- ▶ the Kernel's size,
- ▶ the Stride,
- ▶ the Padding.


The size out the output is :

$$d_{\text{out}} = \frac{d_{\text{in}} + 2 \times \text{Padding} - \text{KernelSize}}{\text{Stride}} + 1$$

CONVOLUTIONAL NEURAL NETWORKS

KERNEL SIZE, PADDING AND STRIDE

Kernel size: 3, Padding: 0, Stride: 1




$$d_{\text{out}} = \frac{d_{\text{in}} + 2 \times \text{Padding} - \text{KernelSize}}{\text{Stride}} + 1$$

CONVOLUTIONAL NEURAL NETWORKS

KERNEL SIZE, PADDING AND STRIDE

Kernel size: 5, Padding: 0, Stride: 1




$$d_{\text{out}} = \frac{d_{\text{in}} + 2 \times \text{Padding} - \text{KernelSize}}{\text{Stride}} + 1$$

CONVOLUTIONAL NEURAL NETWORKS

KERNEL SIZE, PADDING AND STRIDE

Kernel size: 3, Padding: 1, Stride: 1. Padding mode can be 'zeros', 'reflect', 'replicate' or 'circular'.




$$d_{\text{out}} = \frac{d_{\text{in}} + 2 \times \text{Padding} - \text{KernelSize}}{\text{Stride}} + 1$$

CONVOLUTIONAL NEURAL NETWORKS

KERNEL SIZE, PADDING AND STRIDE

Kernel size: 3, Padding: 1, Stride: 2




$$d_{\text{out}} = \frac{d_{\text{in}} + 2 \times \text{Padding} - \text{KernelSize}}{\text{Stride}} + 1$$

CONVOLUTIONAL NEURAL NETWORKS

KERNEL SIZE, PADDING AND STRIDE

Kernel size: 3, Padding: 1, Stride: 3




$$d_{\text{out}} = \frac{d_{\text{in}} + 2 \times \text{Padding} - \text{KernelSize}}{\text{Stride}} + 1$$

CONVOLUTIONAL NEURAL NETWORKS

CNN : CONVOLUTIONAL IN A NETWORK NETWORKS

We can represent a CNN as under this form:




CONVOLUTIONAL NEURAL NETWORKS

CNN : CONVOLUTIONAL IN A NETWORK NETWORKS

Usually, the output of a convolutional block is linear combination of the Convolutional output of every previous channels and a bias:


$$\text{out}_{i,j}(c_{\text{out}}) = \text{bias}(c_{\text{out}}) + \sum_{k=0}^{|c_{\text{in}}|-1} \text{Conv}(\text{input}(k), \text{kernel}_k)_{i,j}$$



CONVOLUTIONAL NEURAL NETWORKS

CNN : CONVOLUTIONAL IN A NETWORK NETWORKS

In practice, we split the image into multiple channels : the three channels RGB to begin with. Then we apply convolutional operation on different scales and then we use a fully connected tail. To change the scale we can use different sub-sampling : Max pooling, Average pooling or Invertible pooling.



CONVOLUTIONAL NEURAL NETWORKS

MAX POOLING

Max Pooling take the maximum within a given sized sub-matrix. In practice, the matrix is size 2×2 in order to reduce the dimension by 4 and doubling the scale.

0.2	1.0	0.3	0.8	0.1	0.8	0.6	0.9
0.7	0.3	0.3	0.5	0.4	0.3	0.3	0.4
0.2	0.6	0.7	0.9	0.9	0.1	0.3	0.5
0.8	0.7	0.1	0.3	0.3	0.6	0.9	0.5
0.7	0.1	0.1	0.2	0.8	0.4	0.9	0.7
0.9	0.1	0.8	0.9	0.6	0.3	0.1	0.9
0.8	0.7	0.2	0.7	0.0	0.6	0.9	0.5
0.9	0.5	0.1	0.2	0.0	0.1	0.1	0.4

Max Pooling
Subsampling

1.0	0.8	0.8	0.9
0.8	0.9	0.9	0.9
0.9	0.9	0.8	0.9
0.9	0.7	0.6	0.9

CONVOLUTIONAL NEURAL NETWORKS

AVERAGE POOLING

The Average pooling takes the average value within the sub-matrix.

0.2	1.0	0.3	0.8	0.1	0.8	0.6	0.9
0.7	0.3	0.3	0.5	0.4	0.3	0.3	0.4
0.2	0.6	0.7	0.9	0.9	0.1	0.3	0.5
0.8	0.7	0.1	0.3	0.3	0.6	0.9	0.5
0.7	0.1	0.1	0.2	0.8	0.4	0.9	0.7
0.9	0.1	0.8	0.9	0.6	0.3	0.1	0.9
0.8	0.7	0.2	0.7	0.0	0.6	0.9	0.5
0.9	0.5	0.1	0.2	0.0	0.1	0.1	0.4

Average Pooling
Subsampling

0.5	0.5	0.4	0.6
0.6	0.5	0.5	0.5
0.4	0.5	0.5	0.6
0.7	0.3	0.2	0.5

CONVOLUTIONAL NEURAL NETWORKS

INVERTIBLE POOLING

For Invertible Networks, we can use Invertible Pooling, aka Squeeze. It preserves the information contained in the channels and keeps the dimension constant.

0.2	1.0	0.3	0.8	0.1	0.8	0.6	0.9
0.7	0.3	0.3	0.5	0.4	0.3	0.3	0.4
0.2	0.6	0.7	0.9	0.9	0.1	0.3	0.5
0.8	0.7	0.1	0.3	0.3	0.6	0.9	0.5
0.7	0.1	0.1	0.2	0.8	0.4	0.9	0.7
0.9	0.1	0.8	0.9	0.6	0.3	0.1	0.9
0.8	0.7	0.2	0.7	0.0	0.6	0.9	0.5
0.9	0.5	0.1	0.2	0.0	0.1	0.1	0.4

Invertible Pooling
→
Subsampling

0.2	0.3	0.1	0.6	1.0	0.8	0.8	0.9
0.2	0.7	0.9	0.3	0.6	0.9	0.1	0.5
0.7	0.1	0.8	0.9	0.1	0.2	0.4	0.7
0.8	0.2	0.0	0.9	0.7	0.7	0.6	0.5
0.7	0.3	0.4	0.3	0.3	0.5	0.3	0.4
0.8	0.1	0.3	0.9	0.7	0.3	0.6	0.5
0.9	0.8	0.6	0.1	0.1	0.9	0.3	0.9
0.9	0.1	0.0	0.1	0.5	0.2	0.1	0.4

CONVOLUTIONAL NEURAL NETWORKS

CNN : CONVOLUTIONAL IN A NETWORK NETWORKS

Convolutional Neural Networks are more suitable for image processing compared to fully connected networks due to their ability to efficiently handle the spatial relationships between pixels in an image. This is achieved through the use of convolutional layers that apply filters to small portions of an image, rather than fully connected layers that process the entire image as a single vector. Additionally, the shared weights in convolutional layers allow for learning of hierarchical features, reducing the number of parameters in the network and increasing its ability to generalize to new images.

CONVOLUTIONAL NEURAL NETWORKS

CNN IN PRACTICE: CIFAR 10

- ▶ Input shape : $3 \times 32 \times 32$.
- ▶ Number of Classes : 10.
- ▶ Number of training samples (x, y) : 50000.
- ▶ Number of evaluating samples: 10000.



CONVOLUTIONAL NEURAL NETWORKS

CNN IN PRACTICE: CIFAR 10

We will compare three different models:


- ▶ Model 1 : Fully Connected Neural Network with 3.4 million parameters.
- ▶ Model 2 : CNN with 62 thousand parameters.
- ▶ Model 3 : Wider and longer CNN with 5.8 million parameters.

CONVOLUTIONAL NEURAL NETWORKS

MODEL 1


The Net is composed of 4 linear layers with ReLU activations:

- ▶ Linear 3072 $\mapsto$ 1024 + ReLU
- ▶ Linear 1024 $\mapsto$ 256 + ReLU
- ▶ Linear 256 $\mapsto$ 64 + ReLU
- ▶ Linear 64 $\mapsto$ 10 + SoftMax




CONVOLUTIONAL NEURAL NETWORKS

MODEL 1



CONVOLUTIONAL NEURAL NETWORKS

MODEL 1




CONVOLUTIONAL NEURAL NETWORKS

MODEL 2


The Net is composed 2 convolutional layers and 2 linear layers:

- Conv $3 \times 32 \times 32 \mapsto 6 \times 28 \times 28 + \text{ReLU}$
- Max Pooling $6 \times 28 \times 28 \mapsto 6 \times 14 \times 14$
- Conv $6 \times 14 \times 14 \mapsto 16 \times 10 \times 10 + \text{ReLU}$
- Max Pooling $16 \times 10 \times 10 \mapsto 16 \times 5 \times 5$
- Linear $400 \mapsto 120 + \text{ReLU}$
- Linear $120 \mapsto 84 + \text{ReLU}$
- Linear $84 \mapsto 10 + \text{SoftMax}$




CONVOLUTIONAL NEURAL NETWORKS

MODEL 2



CONVOLUTIONAL NEURAL NETWORKS

MODEL 2




CONVOLUTIONAL NEURAL NETWORKS

MODEL 3

The Net is composed 6 convolutional layers and 3 linear layers:

- ▶ Conv $3 \times 32 \times 32 \mapsto 32 \times 32 \times 32$ + BatchNorm2d + ReLU
- ▶ Conv $32 \times 32 \times 32 \mapsto 64 \times 32 \times 32$ + ReLU
- ▶ Max Pooling $64 \times 32 \times 32 \mapsto 64 \times 16 \times 16$
- ▶ Conv $64 \times 16 \times 16 \mapsto 128 \times 16 \times 16$ + BatchNorm2d + ReLU
- ▶ Conv $128 \times 16 \times 16 \mapsto 128 \times 16 \times 16$ + ReLU
- ▶ Max Pooling $128 \times 16 \times 16 \mapsto 128 \times 8 \times 8$
- ▶ Conv $128 \times 8 \times 8 \mapsto 256 \times 8 \times 8$ + BatchNorm2d + ReLU
- ▶ Conv $256 \times 8 \times 8 \mapsto 256 \times 8 \times 8$ + ReLU
- ▶ Max Pooling $256 \times 8 \times 8 \mapsto 256 \times 4 \times 4$ + Dropout $p = 0.05$
- ▶ Linear $4096 \mapsto 1024$ + ReLU
- ▶ Linear $1024 \mapsto 512$ + ReLU + Dropout $p = 0.05$
- ▶ Linear $512 \mapsto 10$ + SoftMax


We have added Batch Normalization to improve the training stability and Drop Out to reduce overfitting.



CONVOLUTIONAL NEURAL NETWORKS

DROP OUT

Dropout is a regularization technique in neural networks where **during training**, a portion of the nodes are randomly "dropped out" or ignored during each iteration. This helps prevent over-fitting by preventing the model from relying too heavily on any one node. The result is a more robust and generalizable model that can better handle unseen data.



CONVOLUTIONAL NEURAL NETWORKS

BATCH NORMALIZATION

Batch normalization is a technique in deep learning that is used to normalize the activations of a layer within a batch of data. This helps to prevent the problem of vanishing or exploding gradients and also speeds up the training process. By normalizing the activations, batch normalization helps to stabilize the distribution of the inputs to each layer, reducing the covariate shift and allowing the network to learn more effectively.

CONVOLUTIONAL NEURAL NETWORKS

BATCH NORMALIZATION

- 1: **for** each x_i in a mini-batch B of size b **do**
- 2: Compute the mean μ_B and variance σ_B^2 of the features in the mini-batch B .

$$\mu_B = \frac{1}{b} \sum_i x_i \quad \text{and} \quad \sigma_B^2 = \frac{1}{m} \sum_i (x_i - \mu_B)^2$$

- 3: Normalize each feature x_i in the mini-batch B using μ_B and σ_B^2 .

$$\bar{x}_i = \frac{x_i - \mu_B}{\sqrt{\sigma_B^2 + \varepsilon}}$$

- 4: Scale and shift each normalized feature x_i using two learnable parameters γ and β respectively.


$$y_i = \gamma \bar{x}_i + \beta$$

- 5: **end for**

Algorithm 1: Batch Normalization


CONVOLUTIONAL NEURAL NETWORKS

MODEL 3




CONVOLUTIONAL NEURAL NETWORKS

MODEL 3



CONVOLUTIONAL NEURAL NETWORKS


CNN IN PRACTICE: CIFAR 10



CONVOLUTIONAL NEURAL NETWORKS


INTUITION BEHIND CHANNELS

To examine the information captured by different channels in a Neural Network, we can compare their output on a dataset. For a given input x , we can compute the similarity between the output of a specific channel and the same channel for other images in the dataset.



CONVOLUTIONAL NEURAL NETWORKS


INTUITION BEHIND CHANNELS




CONVOLUTIONAL NEURAL NETWORKS

INTUITION BEHIND CHANNELS


Channel 12




Original Image




Best Match



Best Match




Best Match



Best Match


Channel 15




Original Image




Best Match



Best Match




Best Match



Best Match


Channel 16




Original Image




Best Match



Best Match




Best Match



Best Match

Channel 20



Original Image



Best Match



Best Match




Best Match



Best Match


CONVOLUTIONAL NEURAL NETWORKS

INTUITION BEHIND CHANNELS



CONVOLUTIONAL NEURAL NETWORKS

INTUITION BEHIND CHANNELS




RECURRENT NEURAL NETWORKS

Recurrent Networks (RNNs) are a type of neural network that are specifically designed to handle sequential data, whereas CNNs are more suited for image and grid-like data. The main difference between RNNs and CNNs lies in the way they process data, with RNNs considering the sequence of elements and their interdependencies, while CNNs focus on capturing local patterns within the input.

RECURRENT NEURAL NETWORKS

RECURRENT BLOCK

A Recurrent Network is a type of neural network that contains a loop mechanism, allowing previous outputs to be used as inputs for future computations. This creates a form of memory that allows the network to process sequential data with variable-length sequences.




RECURRENT NEURAL NETWORKS

RECURRENT BLOCK

Some of the limitations of Vanilla RNNs:


- ▶ Vanishing gradient problem: The gradient signals used to update the weights during training can become very small, making it difficult to train RNNs effectively.
- ▶ Exploding gradient problem: On the other hand, gradients can become too large and cause numeric instability, making it difficult to train RNNs effectively.
- ▶ Short-term memory: Vanilla RNNs have difficulty retaining information over long periods of time, making them unsuitable for tasks that require remembering information from previous time steps.
- ▶ Computational limitations: RNNs can be computationally intensive, making it difficult to apply them to large sequences of data.
- ▶ Difficulty with parallelization: The sequential nature of RNNs can make it difficult to take advantage of parallel processing to speed up training and inference.



RECURRENT NEURAL NETWORKS

LSTM


Long Short-Term Memory (LSTM) networks are a variant of recurrent neural networks (RNNs) that overcome some of the limitations of traditional RNNs, such as the vanishing gradient problem and difficulty in learning long-term dependencies. LSTM networks introduce memory cells, gates, and a process for updating cells, which allows them to selectively preserve information from previous time steps.



RECURRENT NEURAL NETWORKS

LSTM


Long Short-Term Memory (LSTM) networks are a variant of recurrent neural networks (RNNs) that overcome some of the limitations of traditional RNNs, such as the problem of vanishing gradients and the difficulty of learning long-term dependencies. LSTM networks introduce memory cells, gates, and a process for updating cells, which allows them to selectively preserve information from previous time steps.



RECURRENT NEURAL NETWORKS

LSTM


Long Short-Term Memory (LSTM) networks are a variant of recurrent neural networks (RNNs) that overcome some of the limitations of traditional RNNs, such as the vanishing gradient problem and difficulty in learning long-term dependencies. LSTM networks introduce memory cells, gates, and a process for updating the cells, which allows them to selectively preserve information from previous time steps.



RECURRENT NEURAL NETWORKS

LSTM

Long Short-Term Memory (LSTM) networks are a variant of recurrent neural networks (RNNs) that overcome some of the limitations of traditional RNNs, such as the vanishing gradient problem and difficulty in learning long-term dependencies. LSTM networks introduce memory cells, gates, and a process for updating the cells, which allows them to selectively preserve information from previous time steps.



RECURRENT NEURAL NETWORKS

LIMITS OF LSTM


Limitations of LSTM RNNs:

- ▶ High computational cost: LSTMs are computationally more expensive compared to other traditional neural network models due to the presence of multiple gates and their sequential processing nature.
- ▶ Vanishing Gradient Problem: LSTMs, like any other RNNs, are prone to the vanishing gradient problem when the sequences are too long, making it difficult for the model to learn long-term dependencies.
- ▶ Overfitting: LSTMs are complex models and are more susceptible to overfitting compared to simple feedforward networks.
- ▶ Difficult to parallelize: Due to the sequential nature of LSTMs, they are difficult to parallelize and can take longer to train.
- ▶ Gradient Explosion: LSTMs can also suffer from the gradient explosion problem, where the gradients can become too large and cause numerical instability during training.

RECURRENT NEURAL NETWORKS

GRU


GRU blocks, or Gated Recurrent Units, are a type of recurrent neural network architecture that are similar to LSTMs in their function and ability to process sequential data. GRUs were introduced as a simplification of LSTMs, with the aim of reducing the number of parameters in the network and improving computational efficiency. GRUs achieve this by merging the forget and input gates in LSTMs into a single update gate, effectively combining the two operations in a single step.



RECURRENT NEURAL NETWORKS

GRU


GRU blocks, or Gated Recurrent Units, are a type of recurrent neural network architecture that are similar to LSTMs in their function and ability to process sequential data. GRUs were introduced as a simplification of LSTMs, with the aim of reducing the number of parameters in the network and improving computational efficiency. GRUs achieve this by merging the forget and input gates in LSTMs into a single update gate, effectively combining the two operations in a single step.



RECURRENT NEURAL NETWORKS

GRU

GRU blocks, or Gated Recurrent Units, are a type of recurrent neural network architecture that are similar to LSTMs in their function and ability to process sequential data. GRUs were introduced as a simplification of LSTMs, with the aim of reducing the number of parameters in the network and improving computational efficiency. GRUs achieve this by merging the forget and input gates in LSTMs into a single update gate, effectively combining the two operations in a single step.



RECURRENT NEURAL NETWORKS

LSTM AND GRU

Limitations of GRU RNNs:

- ▶ Computational complexity: GRUs are more computationally efficient than LSTMs but still more complex than feedforward neural networks.
- ▶ Long-term dependencies: GRUs may struggle with capturing long-term dependencies in sequences, although they perform better in this regard than vanilla RNNs.
- ▶ Vanishing gradient problem: GRUs can still be affected by the vanishing gradient problem that plagues all RNN models. This problem makes it difficult for the model to learn from long sequences.
- ▶ Non-stationary data: GRUs may struggle with nonstationary data, where the statistical properties of the data change over time.

RECURRENT NEURAL NETWORKS

APPLICATION OF RNNs

Applications of RNNs:

- ▶ Natural language processing (NLP): Using RNNs for text classification, language translation, and text generation.
- ▶ Time-series prediction: Using RNNs to make predictions based on sequential data, such as stock prices and weather patterns.
- ▶ Speech recognition: Using RNNs for speech-to-text conversion.

TRANSFORMER AND ATTENTION MECHANISM

Transformers and Attention Mechanisms are relatively recent developments in the field of deep learning, which have become popular for processing sequential data, such as natural language processing (NLP) tasks. Unlike Recurrent Neural Networks (RNNs) which process sequential data by repeatedly applying the same set of weights to the inputs over time, Transformers and Attention Mechanisms use self-attention mechanisms to dynamically weight the importance of different elements in the sequence. This enables Transformers to better capture the long-range dependencies between elements in the sequence, leading to improved performance on NLP tasks.

TRANSFORMER AND ATTENTION MECHANISM

SELF-ATTENTION MECHANISM

Self-attention mechanism in transformers is a method of calculating the weight of each input token in a sequence with respect to every other token in the same sequence, resulting in a representation of the input sequence in which the most relevant tokens have the highest weight. Mathematically, the self-attention mechanism can be represented as a dot product between the query (Q), key (K) and value (V) matrices, obtained from the input sequence, followed by a softmax activation to obtain the attention scores. These scores are then used to compute a weighted sum of the value matrix to produce the final representation.


$$\text{Attention}(Q, K, V) = \text{Softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \quad \text{where } Q \in \mathbb{R}^{m \times d_k}, K \in \mathbb{R}^{n \times d_k}, V \in \mathbb{R}^{n \times d_v}$$

TRANSFORMER AND ATTENTION MECHANISM

SELF-ATTENTION MECHANISM

- ▶ Query (Q): Represents the query vector, which is used to calculate the attention scores. Intuitively, the query vector represents the token that we are interested in.
- ▶ Key (K): Represents the key vector, which is used to calculate the attention scores. The key vector helps to determine the importance of each token in the input sequence.
- ▶ Value (V): Represents the value vector, which is used to compute the weighted sum of the values. The value vector provides the information that is used to update the representation of the input sequence.


The resulting weighted sum of the values represents the output of the self-attention mechanism, capturing the relationships between different parts of the input sequence.



TRANSFORMER AND ATTENTION MECHANISM

SELF-ATTENTION MECHANISM


Illustration (source: Understanding Attention Mechanism in Transformer Neural Networks):



TRANSFORMER AND ATTENTION MECHANISM

MULTI-HEAD ATTENTION


In Multi-head Attention, the self-attention mechanism is performed multiple times in parallel with different weight matrices, before being concatenated and once again projected, leading to a more robust representation of the input sequence. The intuition behind the three matrices (Q , K , V) remains the same as in self-attention, with Q representing the query, K the key and V the value. Each head performs an attention mechanism on the input sequence, capturing different aspects and dependencies of the data, before being combined to form a more comprehensive representation of the input.



TRANSFORMER AND ATTENTION MECHANISM

VISUALIZING MULTI-HEAD ATTENTION

Visualizing Self-Attention for Image:
[Link](#)




TRANSFORMER AND ATTENTION MECHANISM

TRANSFORMERS MODEL

Transformers are neural network models that use an encoder-decoder architecture. The encoder takes the input sequence and converts it into a continuous hidden representation, which is then passed to the decoder to generate the output sequence. The architecture of the transformer model is designed to allow the model to process the entire sequence in parallel, rather than processing one element at a time like in traditional RNNs.

Training of transformers involves optimizing a loss function that measures the difference between the model predictions and the true outputs. This loss function is usually based on the cross entropy between the predicted and true sequences.

The encoder-decoder mechanism is commonly referred to as the seq2seq mechanism.



TRANSFORMER AND ATTENTION MECHANISM

TRANSFORMERS MODEL

More information about transformers and specific model architectures will be covered next semester in the course on Applied Deep Learning.

TP2: BUILD AND USE AN AUTOENCODER

FORMAL INTRODUCTION OF AN AUTOENCODER

Definition

An autoencoder is a type of artificial neural network used to learn efficient codings of unlabeled data. It consists of two main components:


- ▶ An **encoder** function: $encoder(x) : \mathbb{R}^d \rightarrow \mathbb{R}^m$
Maps an input x from the input space $\mathbb{R}^d$ to a hidden representation space $\mathbb{R}^m$.
- ▶ A **decoder** function: $decoder(z) : \mathbb{R}^m \rightarrow \mathbb{R}^d$
Maps the hidden representation z back to the original input space $\mathbb{R}^d$.

Goal

The primary goal of an autoencoder is to learn a representation (encoding) for a set of data, typically for the purpose of dimensionality reduction or feature learning. Through training, the autoencoder learns to compress the data from $\mathbb{R}^d$ to $\mathbb{R}^m$ (where $m < d$) and then reconstruct the data back to $\mathbb{R}^d$ as accurately as possible. This process forces the autoencoder to capture the most important features of the data in the hidden representation z .

TP2: BUILD AND USE AN AUTOENCODER

FORMAL INTRODUCTION OF AN AUTOENCODER



TP2: BUILD AND USE AN AUTOENCODER

AUTOENCODERS AND UNSUPERVISED LEARNING

Unsupervised Learning

Autoencoders are a classic example of **unsupervised learning**. In unsupervised learning, the goal is to learn patterns from unlabelled data. Autoencoders learn to compress and decompress the input data without any explicit labels, aiming to capture the underlying structure of the data.

Objective Function

The learning process of an autoencoder is guided by the minimization of a loss function, typically involving a norm that measures the difference between the input and the reconstructed output. Formally, the objective is to minimize:

$$\min_{\theta} \mathbb{E}_{x \sim P} [l(x, \text{decoder}_{\theta}(\text{encoder}_{\theta}(x)))]$$

where x is the input data, θ represents the parameters of the encoder and decoder, and l is a loss function.

TP2: BUILD AND USE AN AUTOENCODER

AUTOENCODERS AND UNSUPERVISED LEARNING

Unsupervised Learning

Autoencoders are a classic example of **unsupervised learning**. In unsupervised learning, the goal is to learn patterns from unlabelled data. Autoencoders learn to compress and decompress the input data without any explicit labels, aiming to capture the underlying structure of the data.

Objective Function

The learning process of an autoencoder is guided by the minimization of a loss function, typically involving a norm that measures the difference between the input and the reconstructed output. Formally, the objective is to minimize:

$$\min_{\theta} \mathbb{E}_{x \sim P} \left[\|x - \hat{x}\|_2^2 \right]$$


where x is the input data, θ represents the parameters of the encoder and decoder and $\hat{x} = \text{decoder}_{\theta}(\text{encoder}_{\theta}(x))$ in the reconstruction.

APPLICATIONS OF AUTOENCODERS

- ▶ **Reduce the size of the data to transfer:** Autoencoders can compress data into a lower-dimensional space, facilitating faster data transfer by reducing the amount of data that needs to be transmitted.


APPLICATIONS OF AUTOENCODERS

- ▶ **Reduce the size of the data to transfer:** Autoencoders can compress data into a lower-dimensional space, facilitating faster data transfer by reducing the amount of data that needs to be transmitted.
- ▶ **Denoise image:** By learning to ignore the "noise" in the input data, autoencoders can reconstruct cleaner versions of noisy images, effectively removing the noise.




APPLICATIONS OF AUTOENCODERS

- ▶ **Reduce the size of the data to transfer:** Autoencoders can compress data into a lower-dimensional space, facilitating faster data transfer by reducing the amount of data that needs to be transmitted.
- ▶ **Denoise image:** By learning to ignore the "noise" in the input data, autoencoders can reconstruct cleaner versions of noisy images, effectively removing the noise.
- ▶ **Anomaly detection:** Autoencoders can learn the normal patterns within data. Deviations from these patterns, when the reconstruction error is high, can indicate anomalies or outliers in the data.



APPLICATIONS OF AUTOENCODERS

- ▶ **Reduce the size of the data to transfer:** Autoencoders can compress data into a lower-dimensional space, facilitating faster data transfer by reducing the amount of data that needs to be transmitted.
- ▶ **Denoise image:** By learning to ignore the "noise" in the input data, autoencoders can reconstruct cleaner versions of noisy images, effectively removing the noise.
- ▶ **Anomaly detection:** Autoencoders can learn the normal patterns within data. Deviations from these patterns, when the reconstruction error is high, can indicate anomalies or outliers in the data.



TP2: BUILD AND USE AN AUTOENCODER

YOUR TURN !

Get the TP2 on the course website and start working on it.

- ▶ <https://www.alexverine.com>
- ▶ Teaching
- ▶ Deep Learning II
- ▶ Lien Notebooks Python

Part III

TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING


TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING

- ▶ **Batch Normalization** - Seen
- ▶ **Dropout** - Seen
- ▶ **Data Augmentation**
- ▶ **Learning Rate Scheduling**
- ▶ **Early Stopping**
- ▶ **Gradient Clipping**
- ▶ **Weight Initialization**
- ▶ **Regularization**
- ▶ **GPU Acceleration**

TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING

DATA AUGMENTATION


- ▶ Data Augmentation is a technique to increase the diversity of your training set by applying random (but realistic) transformations to the training images.
- ▶ The goal is to train a model that is robust to these transformations.
- ▶ For example, you can randomly rotate, scale, and flip the images in your training set.
- ▶ This helps expose the model to different aspects of the data and reduce overfitting.



TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING

LEARNING RATE SCHEDULING


- ▶ The learning rate is one of the most important hyperparameters to tune for your deep learning model. The learning rate determines how quickly the model learns the optimal weights and how refined the gradient descent process is.
- ▶ If the learning rate is too high, the model may not converge or converge to a higher loss. If it is too low, the model may take too long to train.



TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING

LEARNING RATE SCHEDULING


- ▶ The learning rate is one of the most important hyperparameters to tune for your deep learning model. The learning rate determines how quickly the model learns the optimal weights and how refined the gradient descent process is.
- ▶ If the learning rate is too high, the model may not converge or converge to a higher loss. If it is too low, the model may take too long to train.
- ▶ Learning rate scheduling is a technique to adjust the learning rate during training. For example, you can start with a high learning rate and then decrease it over time.



TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING

EARLY STOPPING


- ▶ Early stopping is a technique to prevent overfitting by stopping the training process when the model's performance on the validation set starts to degrade.
- ▶ The idea is to monitor the validation loss during training and stop training when the validation loss stops decreasing.



TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING

GRADIENT CLIPPING

- ▶ Gradient clipping is a technique to prevent exploding gradients during training.
- ▶ Exploding gradients occur when the gradients of the loss function with respect to the model's parameters are too large.
- ▶ This can cause the model to diverge and fail to learn.
- ▶ Gradient clipping involves scaling the gradients if their norm exceeds a certain threshold.



TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING

WEIGHT INITIALIZATION

- ▶ Weight initialization is a technique to set the initial values of the weights in the model.
- ▶ The initial values of the weights can have a significant impact on the training process and the final performance of the model.
- ▶ If the weights are initialized too small, the model may not learn effectively. If they are initialized too large, the model may not converge.
- ▶ Common weight initialization techniques include Xavier/Glorot initialization and He initialization.

TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING

WEIGHT INITIALIZATION

- ▶ Xavier/Glorot initialization: The weights are initialized from a normal distribution with mean 0 and variance $2/(n_{\text{in}} + n_{\text{out}})$, where n_{in} and n_{out} are the number of input and output units, respectively. It helps prevent the gradients from vanishing or exploding during training by ensuring that the gradients have a similar scale.
- ▶ He initialization: The weights are initialized from a normal distribution with mean 0 and variance $2/n_{\text{in}}$, where n_{in} is the number of input units. It is commonly used for ReLU activation functions.

TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING

REGULARIZATION

- ▶ Regularization is a technique to prevent overfitting by adding a penalty term to the loss function that discourages the model from learning complex patterns that may not generalize well.
- ▶ L1 regularization adds a penalty term to the loss function that is proportional to the absolute value of the weights. It encourages sparsity in the weights.
- ▶ L2 regularization adds a penalty term to the loss function that is proportional to the square of the weights. It encourages the weights to be small.

TECHNIQUES TO IMPROVE DEEP LEARNING TRAINING

GPU ACCELERATION

CPUs and GPUs are very different in terms of architecture and performance. CPUs are more suited for general-purpose computing tasks, while GPUs are optimized for parallel processing of simple operations, making them ideal for deep learning tasks.

- ▶ GPUs are much faster than CPUs for deep learning tasks because they have many more cores and can perform many more operations in parallel.
- ▶ Deep learning frameworks like PyTorch and TensorFlow are designed to take advantage of GPUs to accelerate the training process.
- ▶ Only the forward and backward passes of the model are executed on the GPU. The data loading and preprocessing are still done on the CPU.

Part IV

DEEP LEARNING AND APPLICATIONS

LEARNING TO ACT WITH DEEP REINFORCEMENT LEARNING

DEEP Q-LEARNING

Deep Q-Learning is a reinforcement learning algorithm that utilizes a neural network to approximate the optimal Q function, which is defined as the expected cumulative reward obtained by following a specific policy. The expected reward can be represented mathematically as follows:

$$R(\pi) = \sum_{t \leq T} E_{p^\pi} [\gamma^t r(s_t, a_t)] ,$$

where $r(s_t, a_t)$ is the reward at time step t , γ is the discount factor, T is the final time step and

$$p^\pi(a_0, a_1, s_1, \dots, a_T, s_T) = p(a_0) \prod_{t=1}^T \pi(a_t | s_t) p(s_{t+1} | s_t, a_t) .$$

The Q function, $Q(s, a)$, represents the expected cumulative reward obtained by taking action a in state s :

$$Q^\pi(s, a) = E_{p^\pi} \left[\sum_{t \leq T} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a \right] .$$

The policy, represented by $\pi(a|s)$, is a probability distribution over actions given a state. The optimal Q function, $Q^*(s, a)$, can be found by solving the Bellman equation:

$$Q^*(s, a) = \mathbb{E}[R|s, a] = \mathbb{E}[r + \gamma \max_{a'} Q^\pi(s', a') | s, a] .$$

LEARNING TO ACT WITH DEEP REINFORCEMENT LEARNING

DEEP Q-LEARNING

The loss in Deep Q-Learning method is the difference between the predicted Q-value and the target Q-value, which is the maximum expected reward obtained from the next state:

$$\mathcal{L}(\theta) = E_{s' \sim \pi^*(\cdot|s,a)} \|r + \gamma \max_{a'} Q(s', a', \theta) - Q(s, a, \theta)\|^2.$$

This loss is used to update the parameters of the deep learning model in order to make the predictions more accurate. The target Q-value is typically computed as the reward obtained from taking an action in the current state, plus the maximum expected reward obtainable from the next state:


1. At the state s_t , select the action a_t with best reward using Q_t and store the results;
2. Obtain the new state s_{t+1} from the environment p ;
3. Store (s_t, a_t, s_{t+1}) ;
4. Obtain Q_{t+1} by minimizing $\mathcal{L}$ from a batch recovered from previously stored results.

LEARNING TO ACT WITH DEEP REINFORCEMENT LEARNING

THE MOUSE GAME

- ▶ A **Mouse** has to feed on food (red) and avoid poison (blue).
- ▶ It has a vision range of 2 squares. So it can see the 25 cells around.
- ▶ The reward for a cheese cell is 0.5, while the reward for eating poison is -1 .



On this example, the mouse behaves randomly.



LEARNING TO ACT WITH DEEP REINFORCEMENT LEARNING

MODEL 1



Fully connected network:



LEARNING TO ACT WITH DEEP REINFORCEMENT LEARNING

MODEL 2

Convolutional network:



LEARNING TO ACT WITH DEEP REINFORCEMENT LEARNING



ϵ -GREEDY ALGORITHM

The ϵ -greedy algorithm is a common exploration strategy used in Deep Q learning. Balances exploration, where the agent tries out new actions and collects new data, and exploitation, where the agent uses the information it already has to select the action with the highest expected reward. The algorithm selects a random action with probability ϵ and the action with the highest Q value with probability $1 - \epsilon$. The value of ϵ decreases over time to gradually shift the focus from exploration to exploitation.

LEARNING TO ACT WITH DEEP REINFORCEMENT LEARNING

MODEL 2+ INCORPORATED ϵ -EXPLORATION

Convolutional network + ϵ -greedy Algorithm
during training:



SYNTHETIC DATA GENERATION WITH GENERATIVE ADVERSARIAL NETWORKS

GANs MODELS

Generative Adversarial Networks (GANs) are a type of deep learning architecture composed of two neural networks, the generator and the discriminator, that are trained in an adversarial manner. The generator network is trained to generate fake data that appears similar to real data, while the discriminator network is trained to distinguish between real and fake data. The loss for GANs is defined as a min-max game, where the generator minimizes the loss function, and the discriminator maximizes it. D and G represent the discriminator and generator networks, respectively, and the goal is to find the optimal configuration for D and G such that the generated samples appear indistinguishable from real data.

SYNTHETIC DATA GENERATION WITH GENERATIVE ADVERSARIAL NETWORKS



GANs MODELS

The loss of a (GAN) is defined as a minimax game between the generator and discriminator models. The generator aims to generate samples that are indistinguishable from real samples, while the discriminator aims to distinguish the generated samples from real samples. The loss function for the generator is defined as $-\log(D(G(z)))$, where D is the discriminator model and $G(z)$ is the generator's output for a random noise vector z . The loss function for the discriminator is defined as $\log(D(x)) + \log(1 - D(G(z)))$, where x is a real sample.

$$\min_D \max_G \mathbb{E}_{x \sim p_{\text{real}}} [\log(D(x))] + \mathbb{E}_{z \sim q} [\log(1 - D(G(z)))]$$


SYNTHETIC DATA GENERATION WITH GENERATIVE ADVERSARIAL NETWORKS

MNIST GENERATION




SYNTHETIC DATA GENERATION WITH GENERATIVE ADVERSARIAL NETWORKS

MNIST GENERATION



SYNTHETIC DATA GENERATION WITH GENERATIVE ADVERSARIAL NETWORKS

MNIST GENERATION



SENTIMENT ANALYSIS WITH TRANSFORMERS AND GRU

BERT


BERT (Bidirectional Encoder Representations from Transformers) is a pre-trained language model developed by Google in 2018. It is a transformer-based architecture that uses a masked language modeling task to generate a deep understanding of the contextual relationships between words in a sentence. BERT can be fine-tuned for various NLP tasks such as sentiment classification by adding a classification layer on top of its pre-trained representations. The model has achieved state-of-the-art performance in a wide range of NLP tasks, making it a popular choice for sentiment analysis.

A bidirectional transformer is a type of transformer architecture in natural language processing (NLP) where information from both past and future contexts is taken into consideration when making predictions. This is achieved by processing the input sequence in two directions, starting from the beginning and the end of the sequence, and concatenating the outputs to obtain the final representation. This allows the model to capture the context both in the forward and backward directions, providing a more comprehensive representation of the input sequence.

SENTIMENT ANALYSIS WITH TRANSFORMERS AND GRU

BERT

BERT model consists of multiple transformer encoder blocks, with a self-attention mechanism, a feedforward neural network and layer normalization, stacked on top of each other. It also includes a positional encoding component to capture the relative position of tokens in a sequence, and a segment encoding component to differentiate between different sequences within the same input.



Bert

DEEP LEARNING 2
(pre-trained)

SENTIMENT ANALYSIS WITH TRANSFORMERS AND GRU

TRAINING BERT

Bert is trained on two tasks, the masked language model and the next sentence prediction. In the masked language model task, a portion of the input sequence is masked and the model must predict the original token based on its context. In the next sentence prediction task, the model receives a pair of sentences and must predict whether the second sentence follows the first one in the context of the input text. Both of these tasks are used to train Bert to understand the context of words in a sentence and how they relate to each other, allowing it to perform well on a wide range of natural language processing tasks.

Sentence

There	is	no	curse	in	Elvish,	Entish,	or	the	tongues	of	Men	for	this	treachery.
-------	----	----	-------	----	---------	---------	----	-----	---------	----	-----	-----	------	------------

Basic-level masking

There	is		curse	in	Elvish,		or	the	tongues	of	Men	for	this	
-------	----	--	-------	----	---------	--	----	-----	---------	----	-----	-----	------	--

Entity-level masking

There	is			in	Elvish,	Entish,	or	the				for	this	treachery.
-------	----	--	--	----	---------	---------	----	-----	--	--	--	-----	------	------------


Phrase-level masking

			curse	in	Elvish,	Entish,	or					for	this	treachery.
--	--	--	-------	----	---------	---------	----	--	--	--	--	-----	------	------------

SENTIMENT ANALYSIS WITH TRANSFORMERS AND GRU

SENTIMENT ANALYSIS

Bert can be fine-tuned for sentiment analysis by adding a classifier layer on top of the pretrained Bert model. The layer is trained on a labeled sentiment analysis dataset to predict the sentiment of a given input sequence, which can be a sentence, paragraph, or document. Fine-tuning the model allows it to learn the specific nuances of sentiment in the target task and produce improved results compared to using the pre-trained model alone.




SENTIMENT ANALYSIS




SENTIMENT ANALYSIS WITH TRANSFORMERS AND GRU

SENTIMENT ANALYSIS



SENTIMENT ANALYSIS WITH TRANSFORMERS AND GRU

SENTIMENT ANALYSIS



SENTIMENT ANALYSIS WITH TRANSFORMERS AND GRU

SENTIMENT ANALYSIS

```
In [117]: predict_sentiment(model, tokenizer, "This film is terrible")
```

```
Out[117]: 0.028073227033019066
```

```
In [118]: predict_sentiment(model, tokenizer, "This film is great")
```

```
Out[118]: 0.9749133586883545
```

```
In [119]: predict_sentiment(model, tokenizer, "This lecture was quite boring")
```

```
Out[119]: 0.0537508986890316
```

```
In [120]: predict_sentiment(model, tokenizer, "This lecture was challenging")
```

```
Out[120]: 0.6184227466583252
```


```
In [121]: predict_sentiment(model, tokenizer, "This lecture was dense but interesting")
```

```
Out[121]: 0.7105910181999207
```

DENSITY ESTIMATION WITH NORMALIZING FLOWS

ESTIMATING DENSITY

In Machine Learning, the task of *density estimation* consists in estimating the probability density function of a random variable from a set of observations. This is a fundamental problem in statistics and machine learning, with applications in a wide range of fields, including anomaly detection, clustering, and generative modeling. There are several methods for density estimation, including parametric models, non-parametric models, and deep learning models. In this section, we will focus on a specific type of deep learning model called Normalizing Flows, which is used for density estimation and generative modeling.



DENSITY ESTIMATION WITH NORMALIZING FLOWS

OVERVIEW

A Normalizing Flow is usually seen as:

- ▶ a *generative model*,
- ▶ a *bijective mapping*,
- ▶ an *invertible neural network*,
- ▶ a *density estimator*.

DENSITY ESTIMATION WITH NORMALIZING FLOWS

OVERVIEW




Figure. A mapping between two probability distributions
Point to point

DENSITY ESTIMATION WITH NORMALIZING FLOWS

OVERVIEW




Figure. A mapping between two probability distributions
Subset to subset

DENSITY ESTIMATION WITH NORMALIZING FLOWS

MATHEMATICAL FRAMEWORK

Normalizing Flow

A Normalizing Flow is a bijective function between a data space $\mathcal{X}$ and a latent space $\mathcal{Z}$, both subset of $\mathbb{R}^d$.

$$\begin{aligned} F : \mathcal{X} &\longmapsto \mathcal{Z} \\ x &\longmapsto z = F(x) \end{aligned}$$

Data and Latent Distributions

In theory, a NF maps a target distribution P , ie the data distribution to a simple latent distribution Q .

Usually, Q is set to be a Normal Gaussian multivariate distribution $\mathcal{N}(0_d, I_d)$. p and q are respectively the probability densities of P and Q .

DENSITY ESTIMATION WITH NORMALIZING FLOWS

HOW DOES IT WORK ?

In practice, the mapping is *not perfect*. P^* induces a distribution Q and similarly, the latent distribution Q induces $\hat{P}$, which is the learned distribution. The forward pass F is called the *Normalizing* direction while the inverse pass F^{-1} is called the *Generative* direction.




Figure. 1D Normalizing Flow process.

DENSITY ESTIMATION WITH NORMALIZING FLOWS

INDUCED PROBABILITIES ?

Change of Variable Formula


For a bijective and continuous function F and a latent distribution Q , the distribution induced by Q and F is defined through the *change of variable formula*:

$$\forall x \in \mathcal{X}, \quad \hat{p}(x) = |\det \text{Jac}_F(x)| q(F(x)). \quad (1)$$

DENSITY ESTIMATION WITH NORMALIZING FLOWS

INDUCED PROBABILITIES ?


$$\forall x \in \mathcal{X}, \quad \hat{p}(x) = |\det \text{Jac}_F(x)| q(F(x)).$$



DENSITY ESTIMATION WITH NORMALIZING FLOWS

INDUCED PROBABILITIES ?


$$\forall x \in \mathcal{X}, \quad \hat{p}(x) = |\det \text{Jac}_F(x)| q(F(x)).$$



DENSITY ESTIMATION WITH NORMALIZING FLOWS

INDUCED PROBABILITIES ?


$$\forall x \in \mathcal{X}, \quad \hat{p}(x) = |\det \text{Jac}_F(x)| q(F(x)).$$



DENSITY ESTIMATION WITH NORMALIZING FLOWS

INDUCED PROBABILITIES ?


$$\forall x \in \mathcal{X}, \quad \hat{p}(x) = |\det \text{Jac}_F(x)| q(F(x)).$$



DENSITY ESTIMATION WITH NORMALIZING FLOWS

INDUCED PROBABILITIES ?

$$\forall x \in \mathcal{X}, \quad \hat{p}(x) = |\det \text{Jac}_F(x)| q(F(x)).$$



DENSITY ESTIMATION WITH NORMALIZING FLOWS

DENSITY ESTIMATION

To perform density estimation:

1. Draw $x \sim P^*$,
2. Compute $F(x)$ and $|\det \text{Jac}_F(x)|$,
3. Compute $\hat{p}(x) = q(F(x))|\det \text{Jac}_F(x)|$.




Figure. 1D Normalizing Process of Density Estimation.

DENSITY ESTIMATION WITH NORMALIZING FLOWS

DATA GENERATION

To perform data generation:

1. Draw $z \sim Q$,
2. Compute $x = F^{-1}(z)$.




Figure. 1D Normalizing Flow process of Generation.

DENSITY ESTIMATION WITH NORMALIZING FLOWS

LOG-LIKELIHOOD

Loss

The objective is to approximate P^* with $\hat{P}$. We can minimize the Kullback-Leiber Divergence :

$$\theta = \arg \min_{\theta} \mathcal{D}_{\text{KL}}(P^* \parallel \hat{P}).$$

This is equivalent to maximizing the log likelihood :

$$\theta = \arg \max_{\theta} \mathbb{E}_{x \sim \mathcal{X}} [\log \hat{p}(x)] .$$

DENSITY ESTIMATION WITH NORMALIZING FLOWS

LOG-LIKELIHOOD

$$\mathcal{D}_{\text{KL}}(P^* \parallel \hat{P}) = \int_{\mathcal{X}} p^*(x) \log \left(\frac{p^*(x)}{\hat{p}(x)} \right) dx$$

$$\text{nll} = -\mathbb{E}_{x \sim \mathcal{X}} [\log \hat{p}(x)]$$

DENSITY ESTIMATION WITH NORMALIZING FLOWS

LEARNING STEPS




Figure. Learning Process for a 1D Normalizing Flow.

IMAGE SEGMENTATION WITH U-NET

IMAGE SEGMENTATION

Image segmentation is the process of partitioning an image into multiple segments or regions based on the characteristics of the pixels. It is a fundamental task in computer vision and has applications in various fields, including medical imaging, autonomous driving, and satellite image analysis. There are several methods for image segmentation, including thresholding, clustering, and deep learning-based approaches. In this section, we will focus on a deep learning model called U-Net, which is commonly used for image segmentation tasks.



IMAGE SEGMENTATION WITH U-NET

APPLICATIONS OF IMAGE SEGMENTATION

Image segmentation has a wide range of applications in computer vision and image processing, including:

- ▶ Medical Imaging: Segmentation of organs, tumors, and other structures in medical images.
- ▶ Autonomous Driving: Segmentation of objects such as cars, pedestrians, and road signs in images captured by autonomous vehicles.
- ▶ Satellite Image Analysis: Segmentation of land cover types, buildings, and other features in satellite images.
- ▶ Object Detection: Segmentation of objects in images to localize and classify them.
- ▶ Image Editing: Segmentation of objects for image editing tasks such as background removal and image compositing.

IMAGE SEGMENTATION WITH U-NET


TYPE OF SEGMENTATIONS

Image segmentation can be broadly classified into two types: semantic segmentation and instance segmentation.


- ▶ **Semantic Segmentation:** Semantic segmentation assigns a class label to each pixel in an image, such as road, car, person, etc. The goal is to partition the image into semantically meaningful regions.
- ▶ **Instance Segmentation:** Instance segmentation goes a step further than semantic segmentation by distinguishing between different instances of the same class. It assigns a unique label to each object instance in the image.



Original Image



Semantic Segmentation



Instance Segmentation

IMAGE SEGMENTATION WITH U-NET

U-NET ARCHITECTURE

U-Net is a convolutional neural network architecture designed for image segmentation tasks. It consists of an encoder-decoder structure with skip connections that allow the model to capture both local and global features in the input image.




IMAGE SEGMENTATION WITH U-NET

LOSS FUNCTIONS

3 main loss functions are used for image segmentation tasks:

- ▶ **Categorical Cross-Entropy Loss:** Used for multi-class segmentation tasks where each pixel is classified into one of multiple classes.
- ▶ **Dice Loss:** A similarity-based loss function that measures the overlap between the predicted segmentation mask and the ground truth mask.
- ▶ **Shaped aware Loss:** A loss function that penalizes the model for making errors near the object boundaries, where the segmentation is most critical.

IMAGE SEGMENTATION WITH U-NET

LOSS FUNCTIONS

3 main loss functions are used for image segmentation tasks:

- ▶ **Categorical Cross-Entropy Loss:** Used for multi-class segmentation tasks where each pixel is classified into one of multiple classes.

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C y_{i,c} \log(\hat{y}_{i,c})$$

- ▶ **Dice Loss:** A similarity-based loss function that measures the overlap between the predicted segmentation mask and the ground truth mask.
- ▶ **Shaped aware Loss:** A loss function that penalizes the model for making errors near the object boundaries, where the segmentation is most critical.

IMAGE SEGMENTATION WITH U-NET

LOSS FUNCTIONS

3 main loss functions are used for image segmentation tasks:

- ▶ **Categorical Cross-Entropy Loss:** Used for multi-class segmentation tasks where each pixel is classified into one of multiple classes.
- ▶ **Dice Loss:** A similarity-based loss function that measures the overlap between the predicted segmentation mask and the ground truth mask.

$$\mathcal{L} = 1 - \frac{2 \sum_{i=1}^N y_i \hat{y}_i}{\sum_{i=1}^N y_i + \sum_{i=1}^N \hat{y}_i}$$

- ▶ **Shaped aware Loss:** A loss function that penalizes the model for making errors near the object boundaries, where the segmentation is most critical.

IMAGE SEGMENTATION WITH U-NET

LOSS FUNCTIONS

3 main loss functions are used for image segmentation tasks:

- ▶ **Categorical Cross-Entropy Loss:** Used for multi-class segmentation tasks where each pixel is classified into one of multiple classes.
- ▶ **Dice Loss:** A similarity-based loss function that measures the overlap between the predicted segmentation mask and the ground truth mask.
- ▶ **Shaped aware Loss:** A loss function that penalizes the model for making errors near the object boundaries, where the segmentation is most critical.

$$\mathcal{L} = -w(x) \left[\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C y_{i,c} \log(\hat{y}_{i,c}) \right]$$

with $w(x)$ a weight function that assigns higher weights to pixels near the object boundaries.

Part V

MANAGING A DEEP LEARNING PROJECT

WHY NOTEBOOKS ARE NOT ENOUGH

WHERE NOTEBOOKS FALL SHORT

Jupyter or Colab Notebooks are great for prototyping, showcasing, and data exploration. However, they are not enough for managing a deep learning project. In particular, they lack the following features:

- ▶ **Reproducibility:** Notebooks are not reproducible. They are not designed to be run multiple times with the same results.
- ▶ **Versioning:** Notebooks are not versioned. It is hard to track changes and revert to versions.
- ▶ **Modularity:** Notebooks are not modular. It is hard to reuse code across different notebooks.
- ▶ **Scalability:** Notebooks are not scalable. They are not designed to run on multiple machines or in the cloud.
- ▶ **Monitoring:** Notebooks are not monitored. It is hard to track the progress of a long-running experiment.
- ▶ **Collaboration:** Notebooks are not collaborative. It is hard to work with multiple people on the same notebook.
- ▶ **Deployment:** Notebooks are not deployable. They are not designed to be run in production.
- ▶ **Performance:** Notebooks are not performant, depending on the amount of log, they can be up to 50% slower than a script.

WHY NOTEBOOKS ARE NOT ENOUGH

WHY SCRIPTS ARE BETTER

Directly running Python script is a better alternative to Jupyter Notebooks for a deep learning project. They allow:

- ▶ **Support for Scheduling and Automation:** Scripts can be scheduled to run at specific times or intervals.
- ▶ **Support for Versioning:** Scripts can be versioned using Git or other version control systems.
- ▶ **Support for Scalability:** Scripts can be run on multiple machines or in the cloud.
- ▶ **Support for Command Line Arguments:** Scripts can take command line arguments for customization.
- ▶ **Support for Ressource Management:** Scripts can be run in the background without blocking the terminal or integrate cluster management tools (e.g. SLURM).
- ▶ **Support for Logging:** Scripts can log information to files for monitoring and debugging.
- ▶ **Support for Modularity:** Scripts can be modularized into functions and classes.

WHY NOTEBOOKS ARE NOT ENOUGH

PROGRAM

In this lecture, we will learn how to manage a deep learning project using Python scripts. We will cover the following topics:

- ▶ **Project Structure:** We will learn how to organize a deep learning project into directories and files for better modularity and reusability.
- ▶ **Dependency Management:** We will learn how to manage dependencies using a virtual environment and a requirements file.
- ▶ **Running code as a script:** We will learn how to run Python code as a script from the command line.
- ▶ **Logging:** We will learn how to log information to files for monitoring and debugging.

MODULARITY


WHY CODE SHOULD BE MODULAR

Writing modular code is essential for managing complex deep learning projects. It offers several advantages:

- ▶ **Reusability:** Modular code allows functions and classes to be reused across different parts of the project, reducing redundancy.
- ▶ **Maintainability:** With a clear structure, modular code is easier to maintain and update, especially in large projects.
- ▶ **Testing and Debugging:** Smaller, independent modules make it easier to test and debug specific parts of the codebase.
- ▶ **Collaboration:** A modular design enables multiple team members to work on different parts of the codebase simultaneously without conflicts.
- ▶ **Scalability:** Modular code can scale more effectively, as each component can be optimized or extended individually.

MODULARITY

ORGANIZING YOUR PROJECT



Explanation:

- ▶ `main.py`: Entry point for running experiments
- ▶ `models.py`, `agents.py`, etc.: Modular files for specific tasks
- ▶ `utils.py`: Utility functions and classes
- ▶ `requirements.txt`: List of dependencies
- ▶ `experiments/`: Separate folder for experiment outputs, logs, and saved models

MODULARITY

KEY PRINCIPLES OF MODULARITY

To build a maintainable and scalable project, follow these modularity principles:

- ▶ **Separation of Concerns:** Each module should have a single responsibility (e.g., separate data processing, model architecture, and training code).
- ▶ **Encapsulation:** Keep functionality within defined boundaries to avoid cross-module dependencies.
- ▶ **Configurable Parameters:** Use configuration files for parameters, ensuring that code can be adapted to new experiments without modification.
- ▶ **Interface Design:** Define clear inputs and outputs for each module, making it easier to swap or update components.

MODULARITY

WRITING MODULAR CODE IN PYTHON

Best practices for maintaining modularity in Python:

- ▶ **Use Functions and Classes:** Encapsulate related functionality in functions and classes for better organization.
- ▶ **Separate Configurations:** Store experiment parameters and settings in YAML or JSON files instead of hardcoding.
- ▶ **Follow a Naming Convention:** Consistent, descriptive names for modules, functions, and variables improve readability.
- ▶ **Documentation:** Document modules with docstrings and comments to clarify each part's role and dependencies.

MODULARITY

DEFINING THE NEURALNET CLASS

Code example:

```
import torch.nn as nn

class NeuralNet(nn.Module):
    def __init__(self):
        super(NeuralNet, self).__init__()
        self.block = None

    def forward(self, x):
        x = self.block(x)
        return x
```

MODULARITY

EXTENDING NEURALNET FOR DIFFERENT ARCHITECTURES

Modularity allows for easy extension of base classes. Here, we extend `NeuralNet` to create specialized architectures:

```
class DenseNet(NeuralNet):
    def __init__(self, input_size, hidden_size, output_size):
        super(DenseNet, self).__init__()
        self.blocks = nn.ModuleList([nn.Linear(input_size, hidden_size),
                                       nn.ReLU(),
                                       nn.Linear(hidden_size, output_size)])

class ConvNet(NeuralNet):
    def __init__(self, input_channels, num_classes):
        super(ConvNet, self).__init__()
        self.conv1 = nn.Conv2d(input_channels=16,
                                kernel_size=3,
                                stride=1,
                                padding=1)

        self.fc = nn.Linear(16 * 28 * 28, num_classes)
        self.blocks = nn.ModuleList([self.conv1, nn.ReLU(), self.fc])
```

LOCAL ENVIRONMENT SETUP

WHY SET UP A LOCAL ENVIRONMENT?

Setting up a local environment is crucial for deep learning projects for several reasons:

- ▶ **Dependency Isolation:** Keeps project-specific dependencies separate from system-wide packages, preventing version conflicts.
- ▶ **Reproducibility:** Ensures consistent dependencies, allowing for reproducible experiments and easier collaboration.
- ▶ **Project Portability:** Makes it simpler to share the project with others, as they can replicate the exact environment.

Tools for Creating Local Environments:

- ▶ `virtualenv`: Lightweight environment creation tool
- ▶ `conda`: Environment management with additional package management capabilities

LOCAL ENVIRONMENT SETUP

WHY USE VIRTUALENV?

`virtualenv` offers several advantages over `conda` for local environments:

- ▶ **Lightweight:** `virtualenv` is faster to create and has a smaller footprint compared to `conda`.
- ▶ **Python Version Agnostic:** Works with different versions of Python, regardless of the system's Python version.
- ▶ **Flexible Integration:** Compatible with `pip` and various Python packaging tools, making it easy to install dependencies from `requirements.txt`.

Creating a Virtual Environment:

```
# Install virtualenv if not already installed
pip install virtualenv

# Create a new virtual environment
virtualenv my_env
```

LOCAL ENVIRONMENT SETUP

ESSENTIAL COMMANDS FOR VIRTUALENV

Once the virtual environment is created, here are the key commands to manage it:

► **Activate the Environment:**

```
# On Windows
my_env\Scripts\activate

# On macOS/Linux
source my_env/bin/activate
```

► **Deactivate the Environment:**

```
deactivate
```

► **Install Dependencies:**

```
# Install a package
pip install <package_name>

# Install from a requirements file
pip install -r requirements.txt
```

Tip: Always activate the environment before installing packages.

LOCAL ENVIRONMENT SETUP

WHY USE CONDA?

conda provides several advantages over `virtualenv` for local environments:

- ▶ **Built-In Package Management:** conda can install packages directly.
- ▶ **Cross-Language Support:** Supports packages beyond Python, such as R or C libraries, making it versatile for data science projects.
- ▶ **Easier Dependency Resolution:** conda handles complex dependencies automatically, reducing conflicts.

Creating a Conda Environment:

```
# Create a new environment
conda create --name my_env python=3.8
```

LOCAL ENVIRONMENT SETUP

ESSENTIAL COMMANDS FOR CONDA

Once the conda environment is created, here are the key commands to manage it:

- ▶ **Activate the Environment:**

```
conda activate my_env
```

- ▶ **Deactivate the Environment:**

```
conda deactivate
```

- ▶ **Install Dependencies:**

```
# Install a specific package  
conda install <package_name>
```

```
# Install from an environment file  
conda env update --file environment.yml
```

Tip: Use `environment.yml` files for reproducible setups when sharing your project.

RUNNING EXPERIMENTS

INTRODUCTION TO ARGPARSE

The `argparse` library in Python allows for flexible command-line argument parsing, which is essential for configuring scripts without modifying code.

- ▶ **Customizable Scripts:** Enable users to specify parameters at runtime.
- ▶ **Improves Code Reusability:** Parameters can be adjusted dynamically, making the script adaptable across tasks.
- ▶ **User-Friendly Interface:** Provides help messages and default values, making it easier to use scripts.

Example:

```
import argparse
parser = argparse.ArgumentParser(description="My script description")
```

RUNNING EXPERIMENTS

ADDING ARGUMENTS

Using `argparse`, we can add different types of arguments with various options:

```
# Adding arguments
parser.add_argument("--epochs", type=int, default=10,
                    help="Number of epochs")
parser.add_argument("--lr", type=float, default=0.001,
                    help="Learning rate")
parser.add_argument("--batch-size", type=int, default=32,
                    help="Batch size")

args = parser.parse_args()
```

Explanation:

- ▶ `-epochs`: Integer argument with a default value of 10.
- ▶ `-lr`: Float argument with a default value of 0.001.
- ▶ `-batch-size`: Integer argument to specify the batch size.

RUNNING EXPERIMENTS

EXAMPLE USAGE OF ARGPARSE

Once arguments are defined, we can access them via `args`:

```
# Accessing arguments
print(f"Train for {args.epochs} epochs with batch size {args.batch_size}")
```

Example of running the script with arguments:

```
python my_script.py
```

```
python my_script.py --epochs 20 --lr 0.01 --batch-size 64
```

```
python my_script.py --epochs 200 --lr 1e-5
```

This allows the script to be configured dynamically without changing the code.

RUNNING EXPERIMENTS

RUNNING PYTHON COMMANDS IN THE TERMINAL (MAC/LINUX)

Running Python scripts from the terminal on Mac and Linux provides flexibility and resource control:

- ▶ **Basic Command:**

```
python my_script.py --epochs 10 --lr 0.001
```

- ▶ **Background Execution:** Running a script in the background frees up the terminal for other tasks.

```
python my_script.py --epochs 10 --lr 0.001 &
```

- ▶ **Monitoring Background Jobs:** Use `jobs` to list background jobs and `fg` to bring them to the foreground.

- ▶ **Importance of Background Execution:** Ideal for long-running tasks, especially on servers or remote machines.

RUNNING EXPERIMENTS

RUNNING PYTHON COMMANDS IN THE TERMINAL (WINDOWS)

Running Python scripts from the Command Prompt or PowerShell on Windows is straightforward:

► **Basic Command:**

```
python my_script.py --epochs 10 --lr 0.001
```

► **Running in Background:**

- Use `pythonw` to launch a new command window:

```
pythonw my_script.py --epochs 10 --lr 0.001
```

- Alternatively, use `start /B` for background execution without opening a new window:

```
start /B python my_script.py --epochs 10 --lr 0.001
```

- **Importance of Background Execution:** Allows other terminal tasks and is useful for long-running processes.

RUNNING EXPERIMENTS

SCHEDULING JOBS WITH CRON AND ALTERNATIVES

Cron is a task scheduler for Unix-based systems that runs scripts or commands at scheduled intervals.

- ▶ **Cron Syntax:** The format is `min hour day month day_of_week command`.
- ▶ **Example:**

```
# Run a script every day at midnight
0 0 * * * /usr/bin/python /path/to/my_script.py --epochs 10
```

Alternative Tools:

- ▶ **MacOS:** Use `launchd`, a built-in tool for scheduling tasks.
- ▶ **Windows:** Use the Task Scheduler, which provides a GUI for managing scheduled tasks.

Why Scheduling Matters: Automating tasks saves time and allows for consistent, timely execution of regular jobs (e.g., retraining a model).

MONITORING EXPERIMENTS

WHY MONITORING MATTERS

Monitoring experiments is essential in deep learning projects for several reasons:

- ▶ **Track Progress:** Provides insights into training dynamics, allowing for better optimization and tuning.
- ▶ **Diagnose Issues Early:** Identify potential issues (e.g., vanishing gradients, overfitting) quickly, saving time and resources.
- ▶ **Reproducibility and Accountability:** Logs and metrics enable reproducibility and allow other collaborators to understand the experiment setup and results.
- ▶ **Model Comparisons:** Facilitates comparison of different runs and hyperparameters to identify the best-performing model.

MONITORING EXPERIMENTS

TOOLS FOR MONITORING EXPERIMENTS

There are various tools for effective monitoring and tracking:

- ▶ **Basic Logging:** Use the `logging` module to track metrics, errors, and messages within code.
- ▶ **JSONL Files:** Store logs as structured data, allowing for detailed tracking and easy parsing.
- ▶ **TensorBoard:** Provides a visual interface for monitoring metrics, loss curves, and other parameters in real-time.
- ▶ **Weights & Biases:** A powerful tool for experiment tracking and visualization, allowing for collaborative analysis and sharing.

These tools support different needs, from lightweight solutions (logging) to more robust experiment tracking (TensorBoard, Weights & Biases).

MONITORING EXPERIMENTS

LOGGING IN PYTHON

The `logging` module in Python is a flexible way to capture and track key information during training:

- ▶ **Purpose:** Logs provide structured, timestamped messages about code execution, helping with debugging and monitoring.
- ▶ **Logging Levels:** `logging` supports different log levels:
 - **INFO:** General messages about code execution and key steps.
 - **WARNING:** Messages for potential issues that are not critical.
 - **ERROR:** Logs errors that may stop part of the execution.
- ▶ **Formats:** Logs can include timestamps, filenames, and message content for better traceability.

MONITORING EXPERIMENTS

SETTING UP A LOGGER

Here's how to set up a basic logger in Python:

```
import logging

# Configure the logger
logging.basicConfig(
    filename='experiment.log',
    level=logging.INFO,
    format='%(asctime)s - %(levelname)s - %(message)s'
)

# Log messages
logging.info("Training started")
logging.warning("Batch size might be too large")
logging.error("Out of memory error")
```

Explanation:

- ▶ filename: Saves logs to a file for permanent record.
- ▶ level: Sets the minimum log level (INFO, WARNING, ERROR) for capture.
- ▶ format: Configures the log format with timestamp, level, and message.

MONITORING EXPERIMENTS

SAVING METRICS AND PLOTS

For effective monitoring, it's important to save, rather than print, key metrics and visualizations:

- ▶ **Metrics Tracking:** Save loss, accuracy, and other key metrics to log files or structured formats (e.g., JSON, JSONL) for later analysis.
- ▶ **Plotting Loss Curves:** Generate and save plots of loss, accuracy, and other metrics during training, rather than printing them in the terminal.
- ▶ **Tools for Visualization:** Use tools like `matplotlib` to save plots, or real-time dashboards such as:
 - **TensorBoard:** Visualize training metrics live for a smoother debugging experience.
 - **Weights & Biases:** Track metrics, visualize experiments, and compare runs.

Tip: Automate saving plots and metrics to ensure nothing is missed during training.

VERSIONING YOUR CODE

WHY GIT MATTERS

Git is essential for managing code in deep learning projects for several reasons:

- ▶ **Version Control:** Git tracks changes over time, enabling easy rollback to previous versions and enhancing reproducibility.
- ▶ **Collaboration:** Git allows multiple contributors to work on the same project, with tools to merge code and resolve conflicts.
- ▶ **Backup and Security:** Code stored on remote servers (e.g., GitHub, GitLab) provides a secure, backed-up version of the project.
- ▶ **Experiment Tracking:** Branching allows you to test experimental features or new ideas without affecting the main code.

VERSIONING YOUR CODE

HOW GIT WORKS

Git is a distributed version control system where every version of the repository is saved locally and remotely (e.g., on GitHub):

- ▶ **Local Repository:** Changes are tracked locally on your machine, allowing for offline work and complete version history.
- ▶ **Remote Repository:** A server (e.g., GitHub, GitLab) stores your code and history, making it accessible for collaboration.
- ▶ **Commit History:** Every change is saved as a "commit," building a history of modifications for easy review.

Each commit is like a snapshot of the project at a specific point in time, allowing for seamless version tracking.

VERSIONING YOUR CODE

ESSENTIAL GIT COMMANDS

Here are some fundamental Git commands to manage your repository:

- ▶ **git pull:** Fetch updates from the remote repository and integrate them into your local repository.

```
git pull origin main
```

- ▶ **git add:** Stage changes for the next commit.

```
git add <file_name>    # Add a specific file
git add .               # Add all changed files
```

- ▶ **git rm:** Remove a file from the repository.

```
git rm <file_name>
```

VERSIONING YOUR CODE

COMMITTING AND PUSHING CHANGES

Committing and pushing changes are key steps in Git:

- ▶ **git commit:** Save a snapshot of the staged changes to the local repository with a message describing the change.

```
git commit -m "Describe the change"
```

- ▶ **git push:** Upload your commits from the local repository to the remote repository.

```
git push origin main
```

Regular commits with clear messages help track progress and make it easier to understand code changes.

VERSIONING YOUR CODE

BRANCHING AND PULL REQUESTS

Git branches enable experimentation and collaboration:

- ▶ **Branches:** Create separate lines of development to test features without affecting the main codebase.

```
git branch new-feature      # Create a branch
git checkout new-feature    # Switch to the branch
```

- ▶ **Merging:** Combine changes from a branch back into the main branch once the feature is ready.


```
git checkout main
git merge new-feature
```

- ▶ **Pull Requests (PRs):** PRs are used on platforms like GitHub to review and approve code before merging it into the main branch.

VERSIONING YOUR CODE

GIT BRANCHING WORKFLOW

The following diagram represents a typical Git workflow with branching, committing, and pushing changes.



REFERENCES I