

Introduction to Pytorch

IPEO Exercise 6

Devis TUIA
Emanuele Dalsasso
Manon Béchaz
Robin Zbinden
EPFL-ECEO

October 17, 2025

In the next three exercises, we will focus on
Deep Learning with Pytorch

- IPEO exercise 6 Introduction to Pytorch with Logistic Regression
- IPEO exercise 7 Image Classification with CNNs
- IPEO exercise 8 Model Training and Regularization

In the next three exercises, we will focus on
Deep Learning with Pytorch

- IPEO exercise 6 Introduction to Pytorch with Logistic Regression
- IPEO exercise 7 Image Classification with CNNs
- IPEO exercise 8 Model Training and Regularization

In the next three exercises, we will focus on
Deep Learning with Pytorch

- IPEO exercise 6 Introduction to Pytorch with Logistic Regression
- IPEO exercise 7 Image Classification with CNNs
- IPEO exercise 8 Model Training and Regularization



EPFL Pytorch is Numpy with gradients



- package for numerical operations in Python
- first released 1995
- basis for several classical machine learning packages, like scikit-learn



- package for deep learning in Python
- first released 2016
- basis for deep learning and de-facto standard library for scientific applications
- **extends numpy with automatic differentiation and gradients**
- **allows working with GPUs**

EPFL Pytorch is Numpy with gradients



- package for numerical operations in Python
- first released 1995
- basis for several classical machine learning packages, like scikit-learn



- package for deep learning in Python
- first released 2016
- basis for deep learning and de-facto standard library for scientific applications
- **extends numpy with automatic differentiation and gradients**
- **allows working with GPUs**

EPFL Pytorch is Numpy with gradients



- package for numerical operations in Python
- first released 1995
- basis for several classical machine learning packages, like scikit-learn



- package for deep learning in Python
- first released 2016
- basis for deep learning and de-facto standard library for scientific applications
- extends numpy with automatic differentiation and gradients
- allows working with GPUs



- package for numerical operations in Python
- first released 1995
- basis for several classical machine learning packages, like scikit-learn



- package for deep learning in Python
- first released 2016
- basis for deep learning and de-facto standard library for scientific applications
- **extends numpy with automatic differentiation and gradients**
- allows working with GPUs

EPFL Pytorch is Numpy with gradients



- package for numerical operations in Python
- first released 1995
- basis for several classical machine learning packages, like scikit-learn



- package for deep learning in Python
- first released 2016
- basis for deep learning and de-facto standard library for scientific applications
- **extends numpy with automatic differentiation and gradients**
- **allows working with GPUs**

Image I of N gray pixels p_i

$$I = [p_1, p_2, p_3, \dots, p_N]$$

```
from torch import tensor
I = tensor([[1],[0.8]...],
           requires_grad=True)
```

EPFL Automatic Differentiation

4

Image \mathbf{I} of N gray pixels p_i

$$\mathbf{I} = [p_1, p_2, p_3, \dots, p_N]$$

some function $\mathbf{I} \rightarrow s$, e.g.,
summation

$$s = \sum_{i=1}^N p_i$$

```
from torch import tensor
I = tensor([[1],[0.8]...],
           requires_grad=True)
```

summation function

```
s = I.sum()
```

EPFL Automatic Differentiation

4

Image \mathbf{I} of N gray pixels p_i

$$\mathbf{I} = [p_1, p_2, p_3, \dots, p_N]$$

some function $\mathbf{I} \rightarrow s$, e.g.,
summation

$$s = \sum_{i=1}^N p_i$$

analytic gradient calculation
“effect of one pixel on the sum”

$$\frac{\partial s}{\partial p_1} = \sum_{i=1}^N p_i \frac{\partial}{\partial p_1} = 1$$

```
from torch import tensor
I = tensor([[1],[0.8]...],
           requires_grad=True)
```

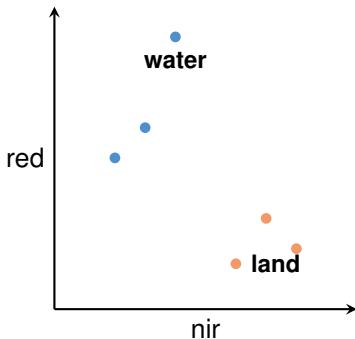
summation function

```
s = I.sum()
```

automatic differentiation

```
s.backward() # backprop
# display grad
I.grad[0]
> 1.
```


Objective: map an input pixel $\mathbf{x} = (x_{\text{green}}, x_{\text{red}}, x_{\text{nir}})$ \rightarrow y probability of “water” $[0,1]$



Objective: map an input pixel $\mathbf{x} = (x_{\text{green}}, x_{\text{red}}, x_{\text{nir}})$ \rightarrow y probability of “water” $[0,1]$

logistic regression model $f_{\mathbf{A},b}(\mathbf{x})$

$$y = f_{\mathbf{A},b}(\mathbf{x}) = \sigma(\underbrace{\mathbf{A}^T \mathbf{x} + b}_{\text{decision plane}})$$



Objective: map an input pixel $\mathbf{x} = (x_{\text{green}}, x_{\text{red}}, x_{\text{nir}})$ \rightarrow y probability of “water” $[0,1]$

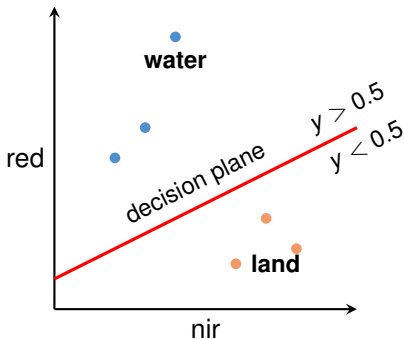
logistic regression model $f_{\mathbf{A},b}(\mathbf{x})$

$$y = f_{\mathbf{A},b}(\mathbf{x}) = \sigma(\underbrace{\mathbf{A}^T \mathbf{x} + b}_{\text{decision plane}})$$

with parameters slope

$\mathbf{A} = (a_{\text{green}}, a_{\text{red}}, a_{\text{nir}})$ and bias b of the decision plane

and a “squishing” sigmoid function $\sigma \mapsto [0, 1]$



How do we find a good decision plane \mathbf{A}, b ?

We solve the optimization objective

$$\mathbf{A}, b = \arg_{\mathbf{A}, b} \min \sum_{i=1}^M \mathcal{L}(f_{\mathbf{A}, b}(\mathbf{x}_i), y_i)$$

in words:

- we find parameters \mathbf{A}, b
- that minimize (arg min) the error $\sum_{i=1}^M \mathcal{L}(f_{\mathbf{A}, b}(\mathbf{x}_i), y_i)$ of
- model output $f_{\mathbf{A}, b}(\mathbf{x}_i)$
- with ground truth y_i over a training dataset of M examples.

we implement arg min with gradient descent in  PyTorch


Idea, we start with random \mathbf{A} , b and update parameters iteratively

$$\mathbf{A} \leftarrow \mathbf{A} - \eta \frac{\partial \mathcal{L}}{\partial \mathbf{A}}$$
$$b \leftarrow b - \eta \frac{\partial \mathcal{L}}{\partial b}.$$

- we can think of the gradient $\frac{\partial \mathcal{L}}{\partial b}$ in words: “what **change** in b **changes** the loss \mathcal{L} ”.
- the learning rate η defines the step size and is an important hyper-parameter.


EPFL Water Classification

8



EPFL Land-water classification








training data

optimize decision plane parameters



EPFL Land-water classification



feature space: green - red - nir




EPFL Land-water classification





all pixels \downarrow training data \rightarrow optimize decision plane parameters

feature space: green - red - nir



feature space: green - red - nir






EPFL Outlook: Why logistic regression?

Logistic regression corresponds to a single linear layer

$$\sigma(\mathbf{Ax} + \mathbf{b})$$



Multi-layer perceptron (neural net with dense layers)

$$\sigma(\mathbf{A}_1\mathbf{x} + \mathbf{b}_1) \circ \sigma(\mathbf{A}_2\mathbf{x} + \mathbf{b}_2)$$

