


AI and Data Analysis



Christian Wolf

The speaker



2000

Dipl.Ing. TU Vienna, Austria



2003

PhD in Computer Science
INSA-Lyon



2004

MCF Télécom Physique Strasbourg



2005

MCF INSA-Lyon





2022-now


Principal Scientist, Team Lead
AI for Robotics

Teaching,
academic research

Research
in a private
laboratory

NAVER is a global ICT company with technology leadership in various fields of business







Seoul city 3D reconstruction



AI for Robotics



1.1 Introduction

1 Introduction

1 Introduction: machine learning, a couple of applications [47]

2 A short history of deep learning [10]

3 The basics of machine learning: fitting and generalization [16]

2 Neural networks and PyTorch

1 Frameworks and Tensors *{+PyTorch}* [24]

2 Simple models (linear regression, logistic regression) *{+PyTorch}* [32]

3 Multi layer models + universal approximation theorem *{+PyTorch}* [32]

4 Train/Validation/Test split; Tensorboard *{+PyTorch}* [14]

5 Gradient Backpropagation and Autograd *{+PyTorch}* [17]

6 Stochastic Gradient Descent and Variants (Adam, RMSProp) [15]

7 Shift invariance and Convolutions *{+PyTorch}* [38]

3 Scaling up: vision, transfer, visualization


1 Computer Vision [39]

2 Semi-supervised, Self-supervised learning [8]


5 GPUs – Software *{+CUDA, +PyTorch}* [14]

(...)

Software development ... and data




To go deeper



Ian Goodfellow, Yoshua Bengio, Aaron Courville, « Deep Learning », MIT Press



PyTorch online tutorials
<https://pytorch.org/>



Christopher Bishop, « Pattern Recognition and Machine Learning », 2006 (*Pre-deeplearning area, but a very pedagogical book on machine learning*)

Learn Python!

For example: <https://learnxinyminutes.com/docs/python/>

Learn X in Y minutes

[Share this page](#)

Select theme: [light](#) [dark](#)

Where X=python

Get the code: [learnpython.py](#)

Python was created by Guido Van Rossum in the early 90s. It is now one of the most popular languages in existence. I fell in love with Python for its syntactic clarity. It's basically executable pseudocode.

Feedback would be highly appreciated! You can reach me at [@louiedinh](#) or louiedinh [at] [google's email service]

Note: This article applies to Python 2.7 specifically, but should be applicable to Python 2.x. Python 2.7 is reaching end of life and will stop being maintained in 2020, it is though recommended to start learning Python with Python 3. For Python 3.x, take a look at the [Python 3 tutorial](#).


It is also possible to write Python code which is compatible with Python 2.7 and 3.x at the same time, using Python [__future__ imports](#). [__future__ imports](#) allow you to write Python 3 code that will run on Python 2, so check out the Python 3 tutorial.

```
# Single line comments start with a number symbol.


""" Multiline strings can be written
    using three "s, and are often used
    as comments
"""


```



Why do we need learning?




The human visual system



Interpretation




Interpretation



Human/hand-crafted
recognition ("Rule-based") is
possible when the data
regularities are simple.


Here: straight lines, some
curves and junctions.





Friend or foe?
Smile or run?

The brain specializes on faces



« compagnon » robot



Figure : Figure-AI "Neo"



Navigation : where am I?
« Visual Landmarks »

Robot navigation




Figure : LAAS-CNRS, Toulouse

Visual SLAM using landmarks

Support for tools (motor control)



Visual servoing





Figure : PRISMA Lab, Université de Naples

Some applications of Deep Learning


Semantic Segmentation



[Fourure, Emonet, Fromont, Muselet, Tremeau, Wolf, BMVC 2017]

PyTorch + computer vision @Tesla

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




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


- (3, 960, 1280) input images
- “ResNet-50 like” dilated backbones
- FPN / DeepLabV3 / UNet -like heads
- ~15 tasks => “prototypes framework”

PyTorch + computer vision @Tesla

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



3D reconstruction: DUST3R



<https://www.youtube.com/watch?v=kl7wCEAFFb0>

[Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, Jerome Revaud, 2023]

Human pose estimation: Anny



<https://www.youtube.com/watch?v=vpAZQrty45Y>

[Romain Brégier, Guénolé Fiche, Laura Bravo-Sánchez, Thomas Lucas, Matthieu Armando, Philippe Weinzaepfel, Gregory Rogez, Fabien Baradel, 2025]

Bullet time



<https://free-view-video.github.io/>

[C. Gao, A. Saraf, J. Kopf, J.-B. Huang, Dynamic View Synthesis from Dynamic Monocular Video, ICCV 2021]


Bullet time in 1999 (Matrix)




Learning inverse problems

Learning is particularly useful in applications, where

1. a forward function is known and easy, but its inverse is difficult and needs to be learned, and
2. Data can be found / generated easily



Super-resolution



[Wei Long, Xingyu Zhou, Leheng Zhang , Shuhang Gu, CVPR 2026]

De-Oldify



<https://github.com/jantic/DeOldify>


Symbolic mathematics

It is easy to: generate random expressions, and derive them.

It is hard to: integrate them

Learn it!

$2 + 3 \times (5 + 2)$, $3x^2 + \cos(2x) - 1$, and $\frac{\partial^2 \psi}{\partial x^2} - \frac{1}{\nu^2} \frac{\partial^2 \psi}{\partial t^2}$:



Learning to control

Task: design a controller for a system.

1. We know the system (equations!):

Solution: design the controller with control theory.


2. We do not know the system but we have expert trajectories (inputs and outputs):

Solution: use Supervised Learning.

3. We know the system (equations!): and do not have any expert trajectories. But given a trajectory, we can know whether the control decisions were satisfactory or not.

Solution: use Reinforcement Learning.

Learning to control



[Tan, Zhang, Coumans, Iscen, Bai, Hafner, Bohez, Vanhouke, RSS 2018]


Learning to control




https://www.youtube.com/watch?v=l44_zbEwz_w

[Boston Dynamics Atlas, 2025]

Taking decisions




{dog, **cat**, avocado, chair, ...}



{0, 1, ... 26, 27, **28**..., 98, 99, ...}



{Left, right, forward, backward, ...}




Motor control

“A blue parrot with a yellow belly
sitting on a branch in a forest”




Decision taking




High dimensional
input data
(e.g. 512x512x3 pixels)


Decision taking: expert knowledge




Decision taking: adding learning




Adding machine learning



Decision taking: deep learning



Decision taking: deep learning



High dimensional
input data
(e.g. 512x512x3 pixels)

Deep Learning:

- Learning from raw signals
- Hierarchical, layered representation
- Different levels of abstraction
- Learning from massive amounts of data, requiring massive compute