LECTURE 6

SELF-SUPERVISED LEARNING

Learning Representations without Labels

PRESS → OR SPACE.

© Marc Rußwurm
Licensed under **Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0)**.
You may share and adapt these slides for teaching, research, and other non-commercial educational purposes with attribution.
Commercial training, paid workshops, consulting seminars, or incorporation into commercial course products require prior permission.

UNIVERSITÄT BONN

Two Questions for Lectures 4–5

LECTURE 4

How does a neural network learn?

LECTURE 5

Why should what it learns work beyond the training data?

Learning Outcomes Roadmap

	Learning outcome	Block
⬜	Explain why representation choice determines task difficulty.	Why Learn Representations?
⬜	Distinguish representation learner vs final decision layer in a network.	Why Learn Representations?

Today: Lecture Structure

Builds on Lecture 4: linear and logistic regression, MLPs, gradient descent and backpropagation.

1

Why Learn Representations?

Start with representation quality and geometric separability.

Why Learn Representations?

Good representations make hard problems easy.

From manual feature design to learned representations.

A difficult task is easy in the right form.

Recap Lecture 1

CHINESE	□□□□□□□□□□□□□□□□□
ARABIC	لقد تعلّم عقلنا أن يفهم هذه الجملة.
HEBREW	המוח שלנו למד להבין את המשפט הזה.
KOREAN	□□□ □□□ □ □□□ □□□□ □□ □□□□□.
LATIN	Our mind has learned to understand this sentence.



Photo by Vitaly Gariev on Pexels

Choosing Suitable Data Representations

Recap Lecture 1

Images

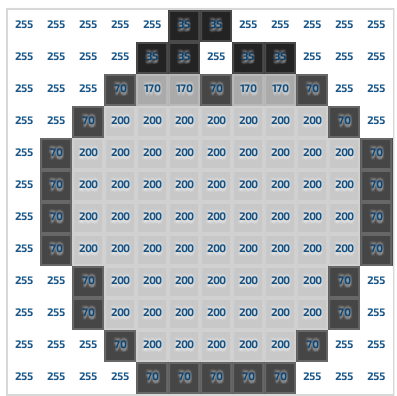

Math

```
XLVII + LXXVIII = ?
```

47 + 78 = 125

Places

COORDINATES

50°43'57.73" N,
7°06'16.63" E

PLACE

Hofgarten, Bonn, Germany

Learning Feature Representations

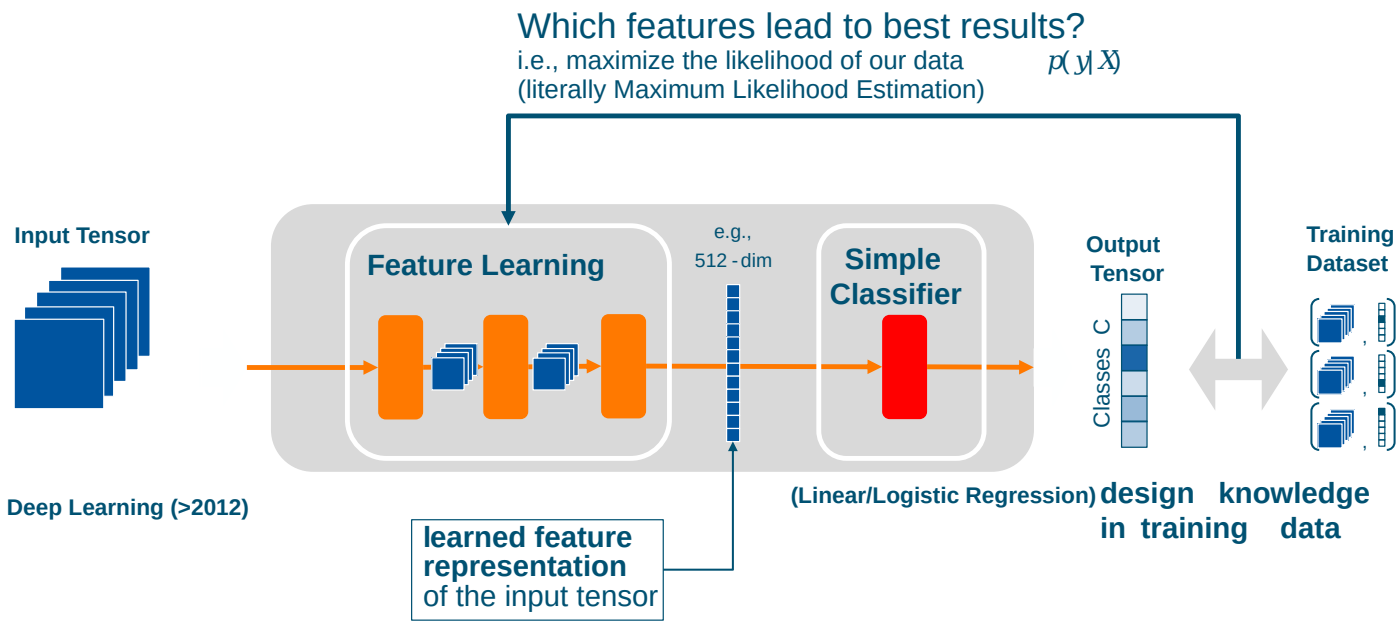

Timeline and Evolution of Deep Learning

<2012	2015 onwards	>2020 →
Classic Machine Learning	**Supervised Deep Learning**	**Self-Supervised Learning**
Feature Design: Which features extract the most information?	**Model Design:** Which architecture is best for my tensors?	**Loss Design:** Which loss function or training strategy?
Data: small labelled	**Data:** large labelled	**Data:** large unlabelled
Models: Random Forest, Boosting, Lin. Regression	**Models:** MLPs, CNNs, RNNs	**Model:** Transformers
Covered in Lectures 3 & 4	*Covered in Lectures 5 & 6*	*Covered from Lecture 7 onwards*

Image Classification

Input Tensor
reflectances

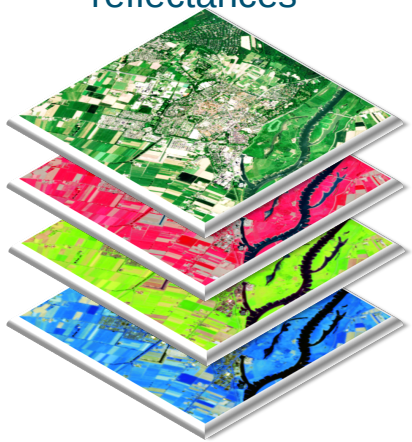

[H x W x D]

Image Classification

(Deep) Model

Output Tensor
class probabilities



[C]

Image Segmentation

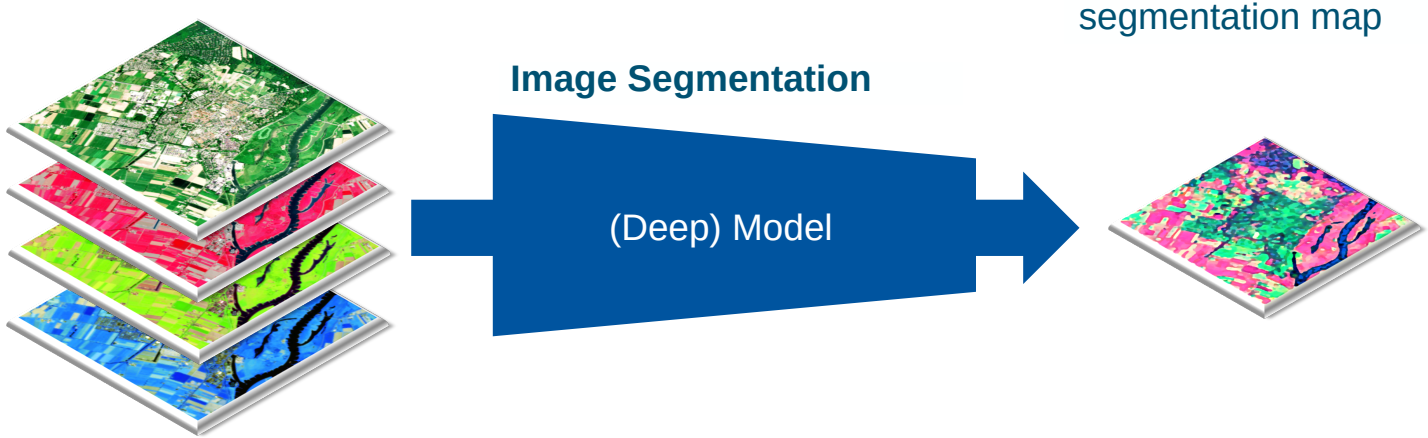

Object Detection

Input Tensor
reflectances



[H x W x 3]

Object Detection

(Deep) Model

Output Tensor
bounding boxes

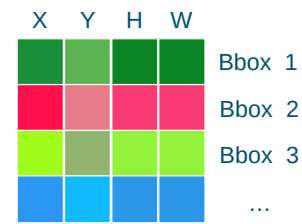

[N x 4]

Time Series Classification

[T x H x W x D]

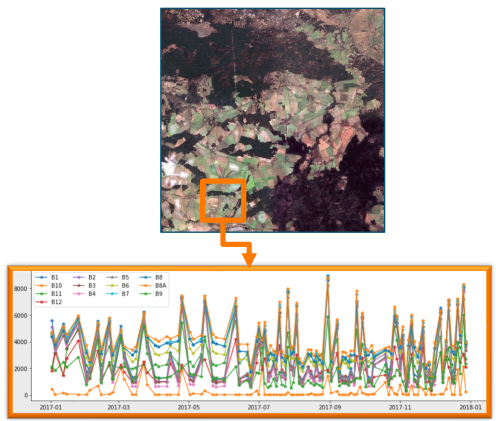

[T x D]

Time Series Classification

(Deep) Model

e.g., classes

[C]

Overview: Deep Model Architectures

Lab 4 Content

1950s | 1990s | 2015 | 2020s

MLP
Linear Projections
Activation Functions
Linear Algebra &
Biological Analogy
Parallelism on GPUs

CNN
Convolutions
Pooling
Recep. Fields

ResNet
Deeper Networks
Batch Normalization
Skip Connections
Classification &
Segmentation Tasks

Transformers
Attention
Mechanism
Self-Attention (SA)
Multi-Head SA
Transformer Layers



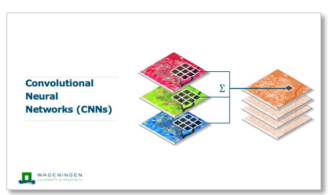



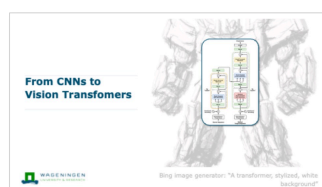

Recap: Learned Representations

- Good representations make difficult problems easier.
- In classic machine learning, we manually select data and engineer features to be a good representation of the problem.
- Deep learning models learn a feature representation that is then well-suited for a final linear classification with logistic regression.

