

Mechatronic Engineering program

Basics of AI and Deep Learning: 10: From Shallow to Deep Learning *In image interpretation...*

Ziemowit Dworakowski
AGH University of Krakow

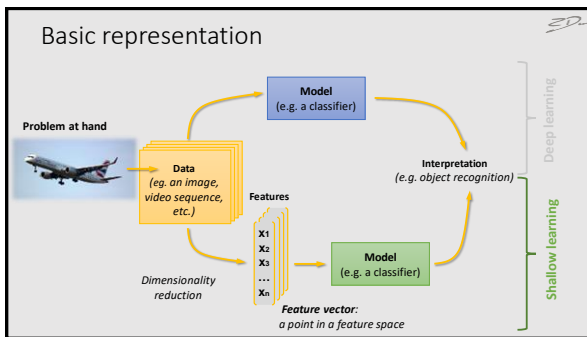
1

AI usage in image processing tasks

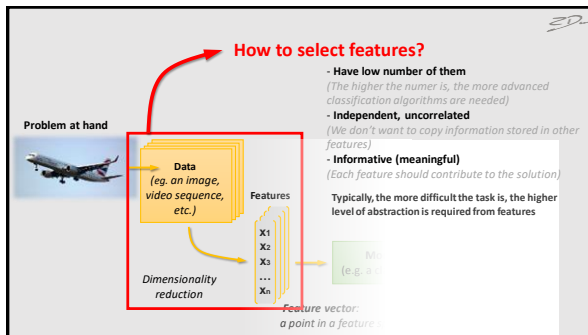
- | | | |
|--|---|--|
| Can be done by
shallow and
deep learning | [| - Classification / Labeling
<i>Putting labels on the image or its parts</i> |
| | | - Intelligent global processing of images
<i>Filter calibration, filtration, image segmentation, enhancing</i> |
| Can be done
by deep
learning only | [| - Modeling – advanced interpretation of vision data
<i>Understanding what is going on in the image (and why)</i> |
| | | - Artificial image generation
<i>Building new images from prompts, augmenting image datasets</i> |

2

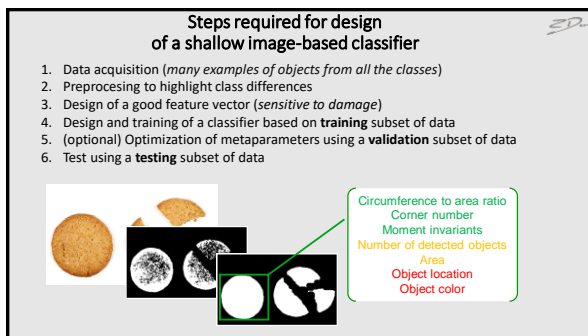
Basic representation



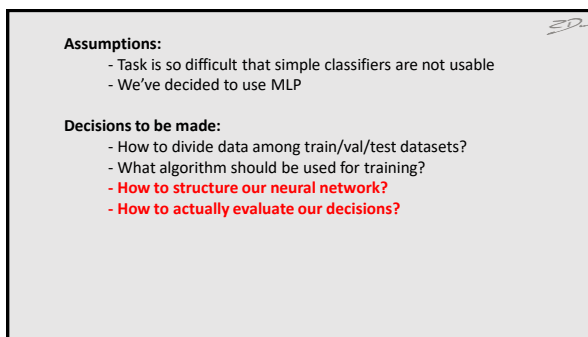
3



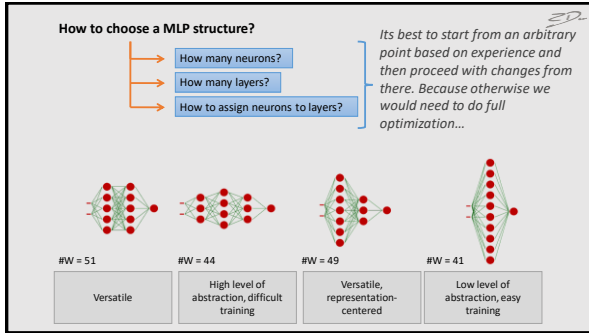
4



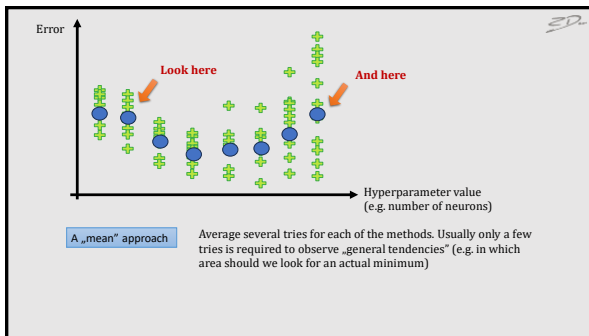
5




6



7



8



9

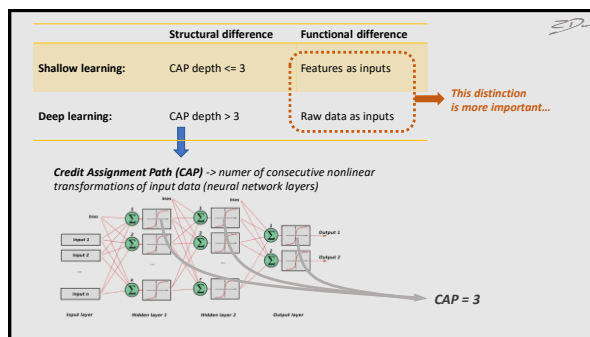
General remarks:

- 1. Classifier should be chosen on a basis of our knowledge of a feature space**
(How many dimensions? How many features? Are the samples clustered?)
The „task type“, (e.g. do we classify cookies or vegetables) is much less relevant.
- 2. Training, testing (and validation) datasets should be separate**
Either we begin with random division of a data into training and testing subsets, or (better!) we gather new portion of data for testing purposes in another experiment.
- 3. Number of degrees of freedom of a classifier (e.g. net weights) should depend on number of data samples**
A good „rule of thumb“ is that for each DOF of a classifier at least 10 data samples are required. If we can't do that, we make sure that overfitting is accounted for!
- 4. Feature quality > Classifier**
Good features allow for easy classification even with a simple classifier. Advanced classifier won't overcome weak features. It is better to spend more time on feature extraction than on classifier configuration.

10

Introduction to deep learning

11

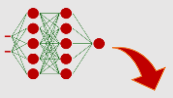



12

How **not** to build your deep learning engine?

We would see:

- Millions of input neurons (for all the pixels of the source)
- Huge amounts of information processed (Millions of neurons in each layer)
- Lack of good feature representations
- Stagnation in training (e.g. due to gradient decay)
- ...

13

So... What actually allows us learning path from raw data to high-level interpretation?

- ReLU and Leaky ReLU activations
- Convolutional kernels
- Pooling
- Regularization
- Fine-tuning

These are basic concepts – most of DL state-of-the-art engines use them in one form or another

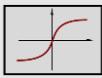
- Adversarial training
- Attention mechanisms
- Autoencoders
- Latent space
- Embeddings
- Reinforcement learning
- Transfer learning
- Transformers
- Recurrent neural networks (including Long-Short-Term-Memory)

These are more advanced and required to understand specialized applications, for instance Midjourney or ChatGPT

14


Sigmoid activation

+ works nice for small nets
- may cause gradient decay in large nets (rendering training to be inefficient)



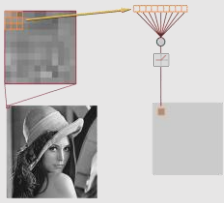
Rectified Linear Unit (ReLU) activation

+ Combats gradient decay in large nets
+ Faster training
-/+ Can cause „neuron death“ problem



15

Convolutional kernel

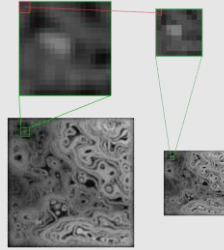


Provides:

- * Convolutional filtration
- * Nonlinear filtration
- * Morphological filtration
- * Extracts low-level features (recognition of „small objects“)
- * Thanks to **weight sharing** one net can do multiple tasks at once
- * We can have **feature maps** as outputs

16

(Max) pooling



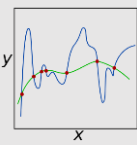
- * We preserve the highest map values
- * We preserve **spatial relations** between features
- * We **decrease dimensionality** of the problem

17

Regularization

In order to prevent memorizing data we can use additional constraints – typically referring to admissible level of task complexity.


In practice we usually **iteratively slightly spoil the classifier**. Whenever generalization is achieved and the classifier begins fitting to noise, regularization factor starts to dominate over parameter-update routine.



For example:

- * In gradient-based training, apart from weight update by the gradient-based policy we also randomly modify all or some weights
- * Some net connections are deleted in-between net training cycles (a „brain damage“ approach)

18



19

Fine tuning

In any case we'll need **a lot of data** and **large computing power**

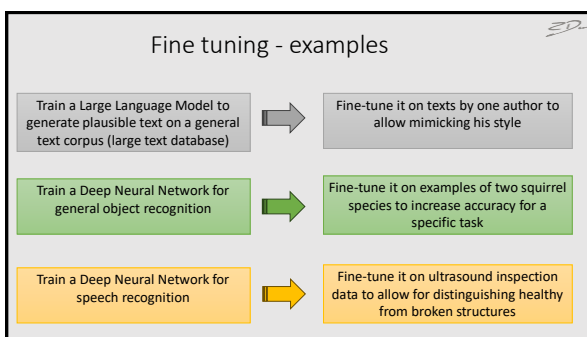
In specialized cases there is often not enough data samples for training
It is often costly and time consuming to train a full model from scratch

Solution:
general training on general data, **fine-tuning** on the most relevant data

↓


We do that generally by reinitializing weights of final layers of the neural network while freezing the remaining layers

20




21

Examples: Matlab DCNN



22

Examples: YOLO v4



23

Summary (topics for test):

- 1) Examples of tasks in intelligent image processing
- 2) Processing path in deep and shallow learning (steps necessary, with example)
- 3) Functional and structural difference between shallow and deep learning
- 4) Differences between various MLP configurations
- 5) How can we optimize MLP structure? (Mean approach, box-plot approach)
- 6) Explain 4 general remarks for classifier training
- 7) Explain sigmoid, ReLU and Leaky ReLU activation functions
- 8) Explain convolutional kernel
- 9) Explain max pooling
- 10) Explain regularization
- 11) Explain fine-tuning (with examples)

(Detailed schemes of DL architectures (like: matlab DCNN or YOLO) will not be necessary for a test)

24