Deep Learning et al.

Architecture and applications

02.11.2022

Mendelson Lab Meeting

What is what?

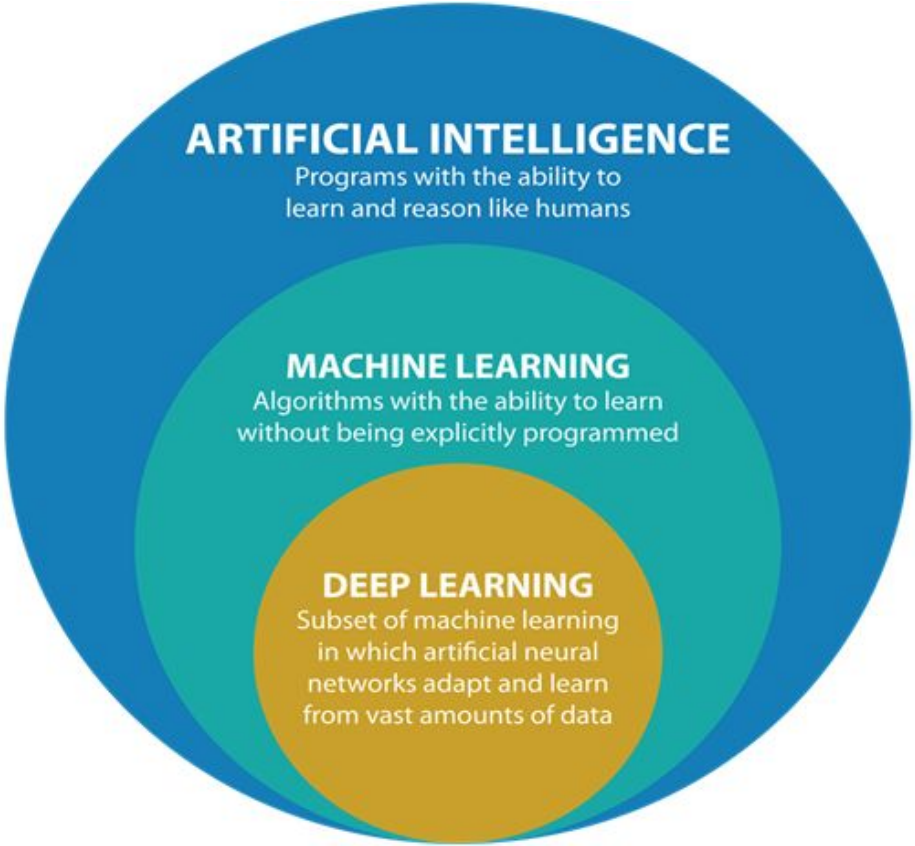
ARTIFICIAL INTELLIGENCE
Programs with the ability to learn and reason like humans
MACHINE LEARNING
Algorithms with the ability to learn without being explicitly programmed
DEEP LEARNING
Subset of machine learning in which artificial neural networks adapt and learn from vast amounts of data

What is what?

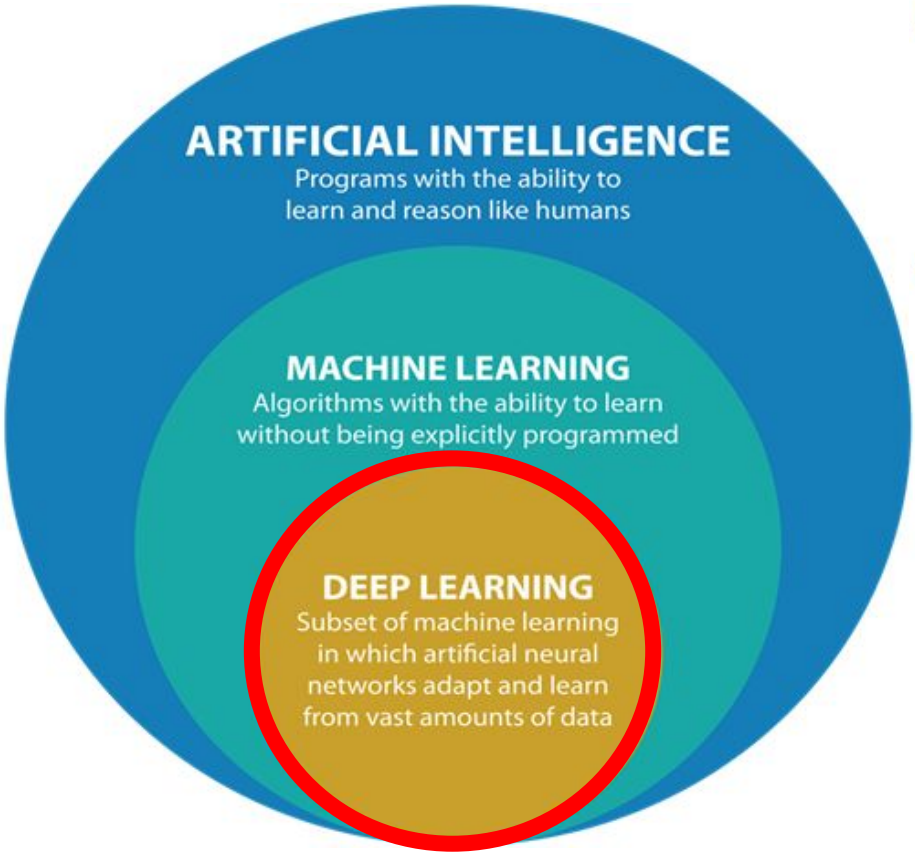
ARTIFICIAL INTELLIGENCE
Programs with the ability to learn and reason like humans
MACHINE LEARNING
Algorithms with the ability to learn without being explicitly programmed
DEEP LEARNING
Subset of machine learning in which artificial neural networks adapt and learn from vast amounts of data

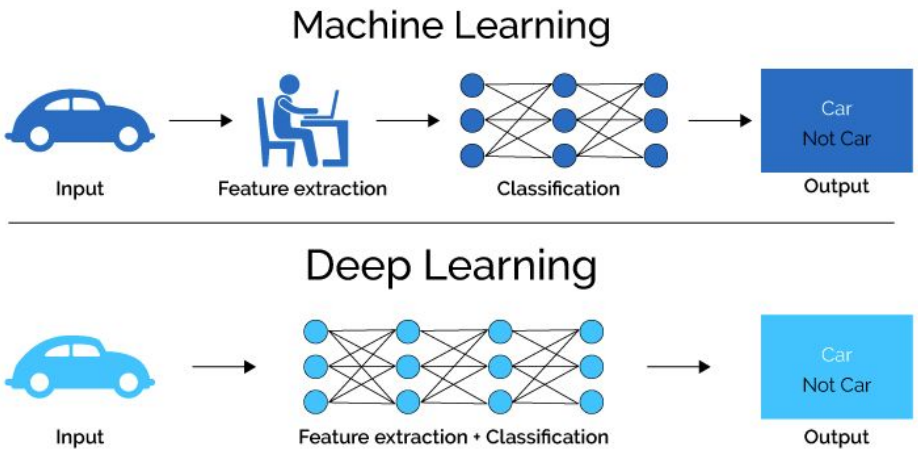
Machine Learning
Input
Feature extraction
Classification
Car
Not Car
Output
Deep Learning
Input
Feature extraction + Classification
Car
Not Car
Output

Deep learning: The different kinds

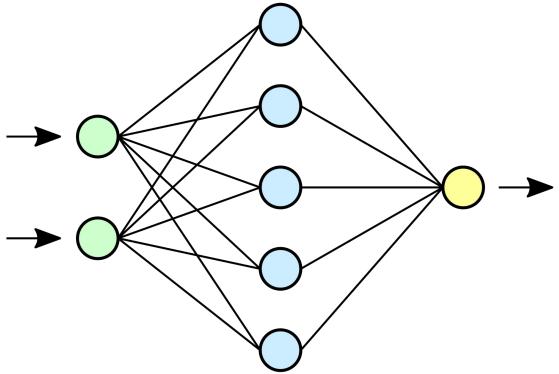

Many acronyms: what do they actually mean?

ANN: artificial NN / **DNN**: deep NN = *generic terms*

CNN or ConvNets: convolutional NN = *most spread type of ANN, feedforward*

GAN: generative adversarial network / **(V)AE**: (variational) autoencoder
= *generative algorithms that involve 2 NNs competing with each other to become more accurate in their predictions*

RNN: recurrent NN = *temporal sequence data (e.g. language →translation)*

ConvNets: Going into the details

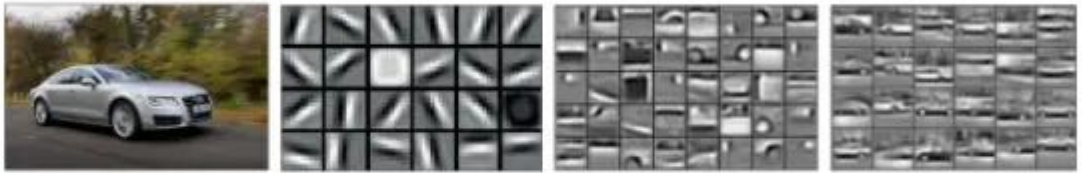

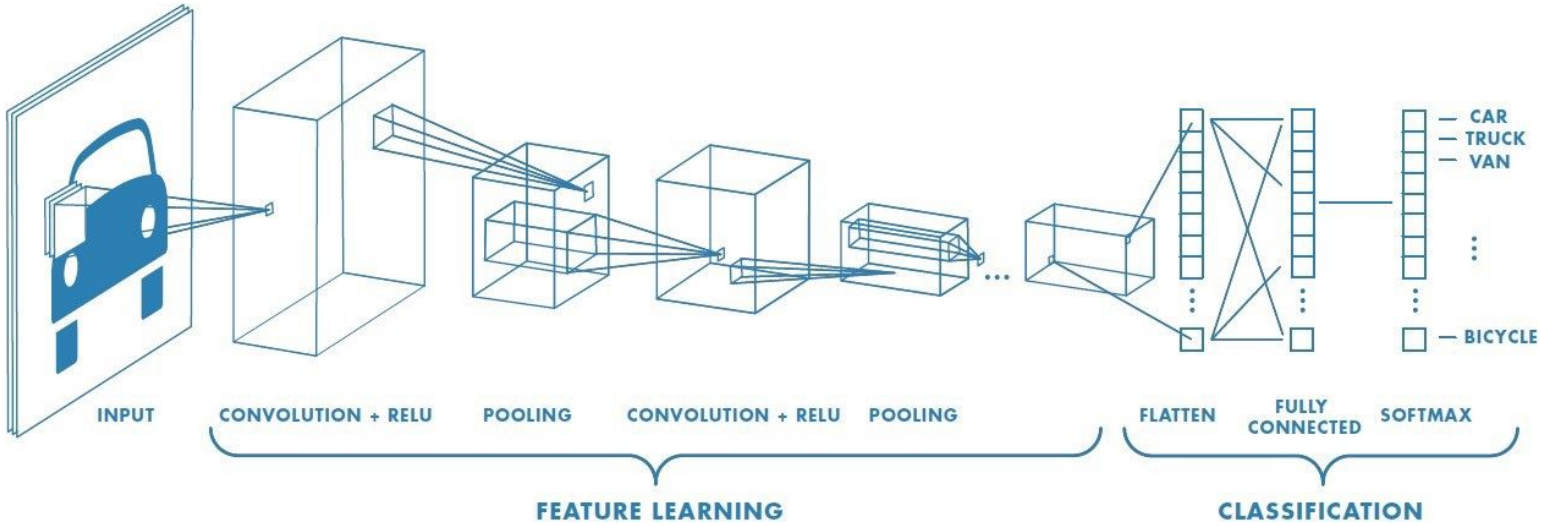

Convolution: .Multiplication of matrices

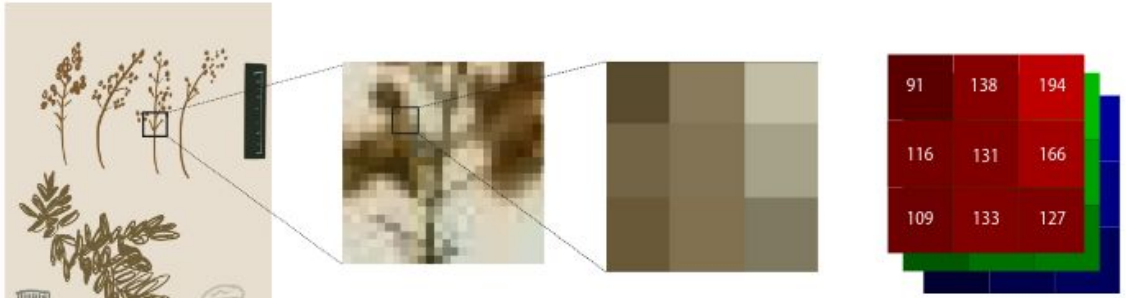

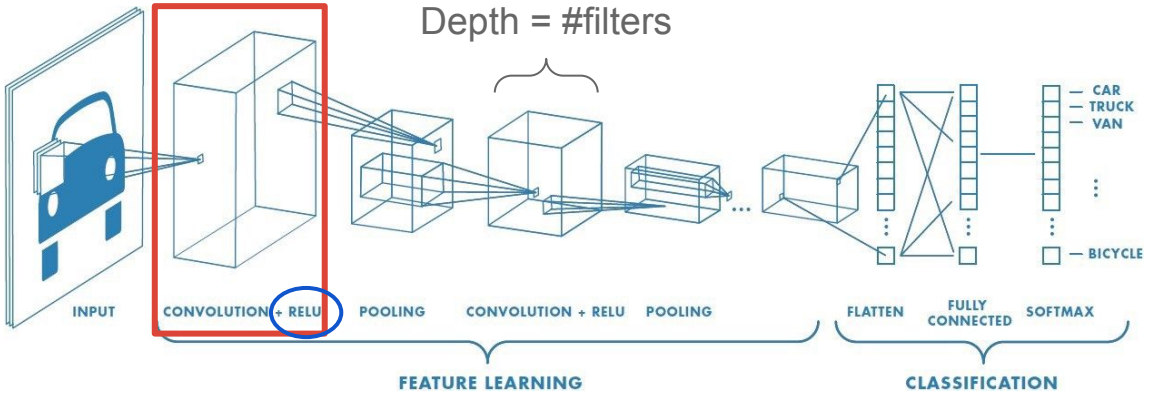

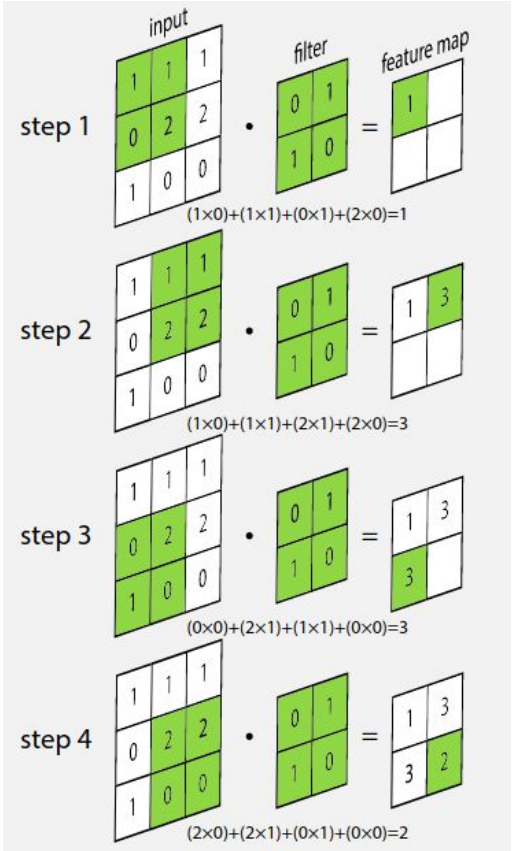

Pooling: Dimension reduction

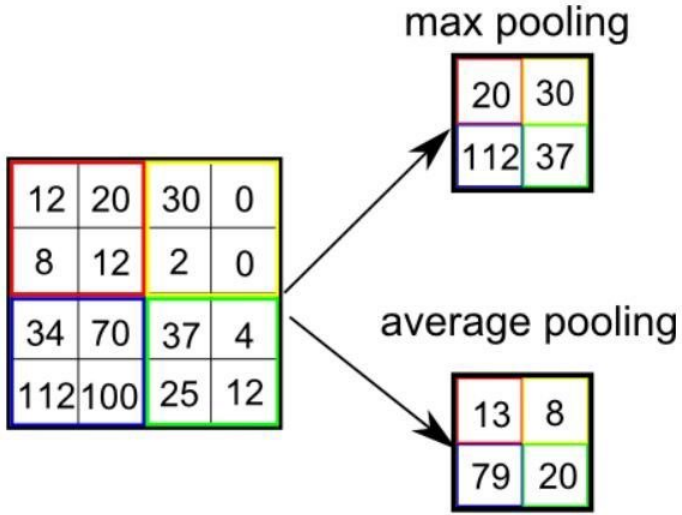

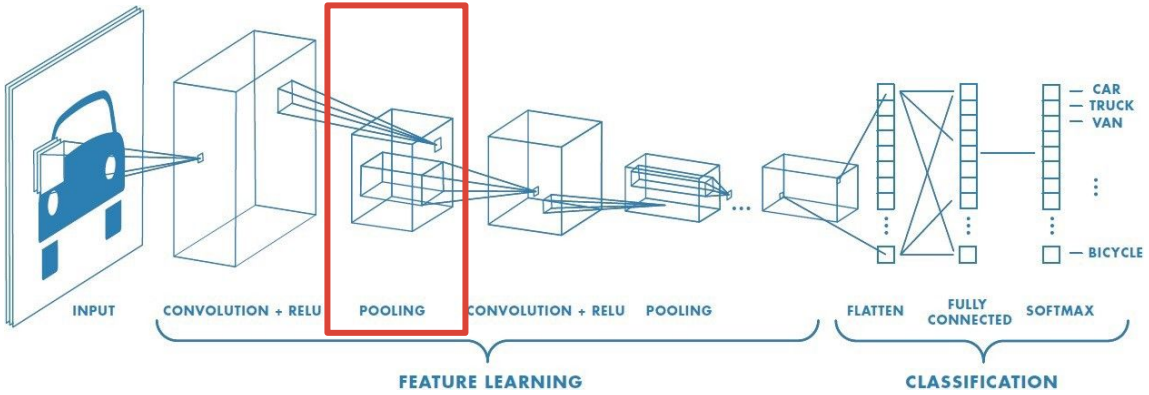

Dense layers / Fully connected layers

Which high level features most strongly correlate to a particular class?

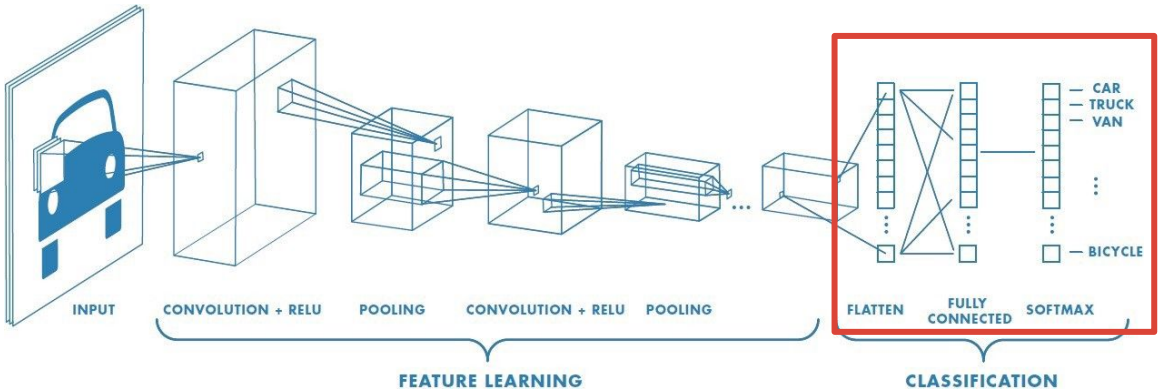

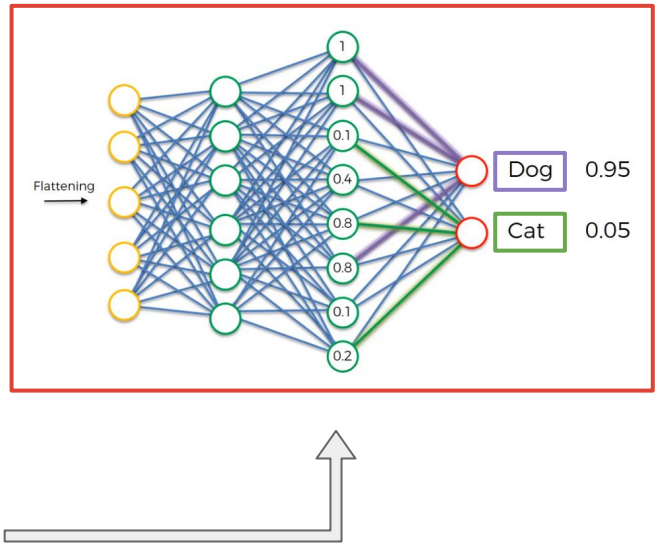

Backpropagation: the training process

1.Forward pass

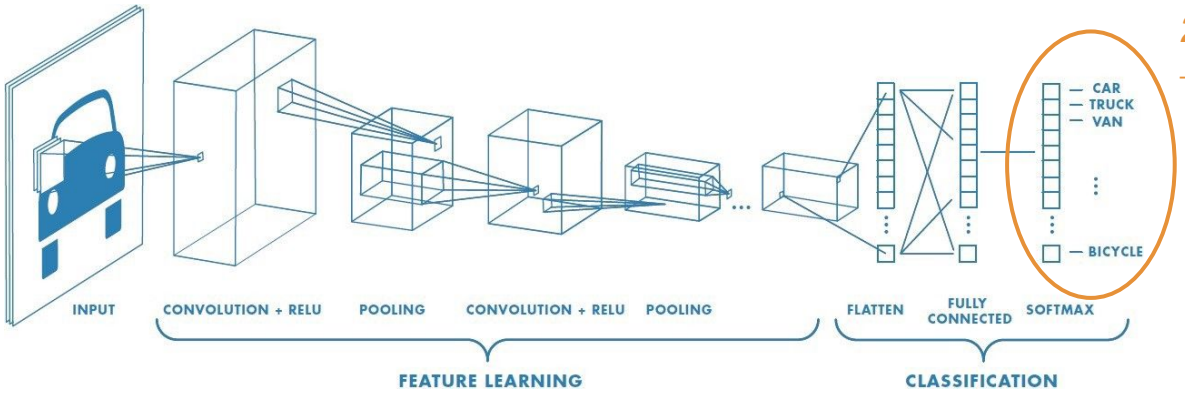

2. Loss function (to minimise)
→ output vs target

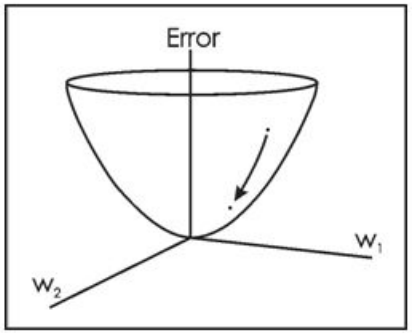

3. Backward pass: which weights contributed the most to the error → adjust them

4. Updated weight = initial weight – learning rate x learning rule (‘gradient’)

Optimisation problem

ss

repeat until convergence {

$$\theta_j := \theta_j - \alpha \frac{\partial}{\partial \theta_j} J(\theta_0, \theta_1)$$

(for $j = 1$ and $j = 0$)

}

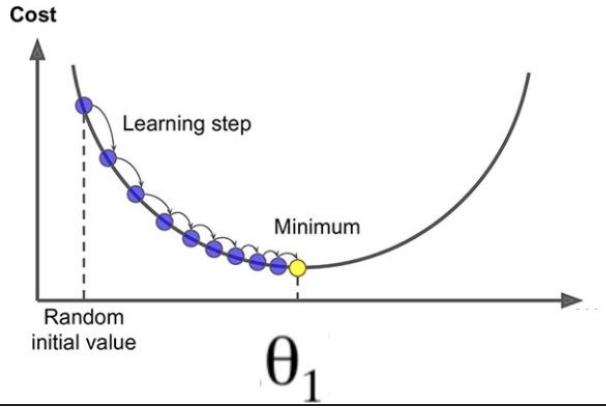

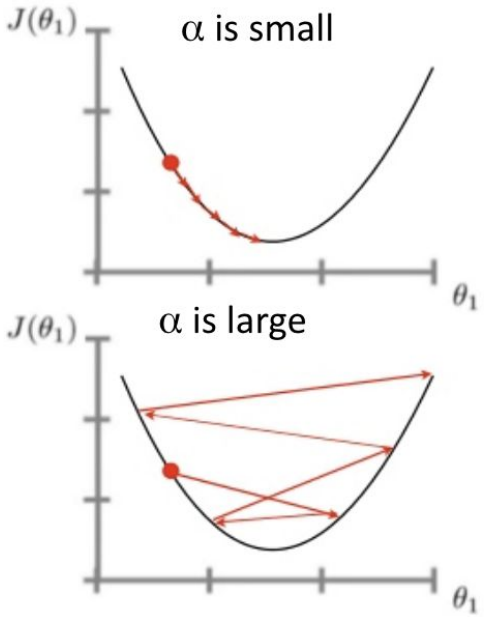

s function (to minimise)
out vs target

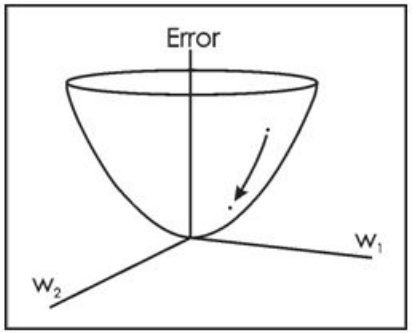

3. Backward pass: which weights contributed the most to the error → adjust them
4. Updated weight = initial weight – learning rate x learning rule ('gradient')

Optimisation problem

Deep learning: What for?

Tasks: classification, object detection, segmentation, individual identification (feature-based)
→ body posture and movement tracking, classification of behaviours
→ genomics (sequence prediction); pop genetics estimations

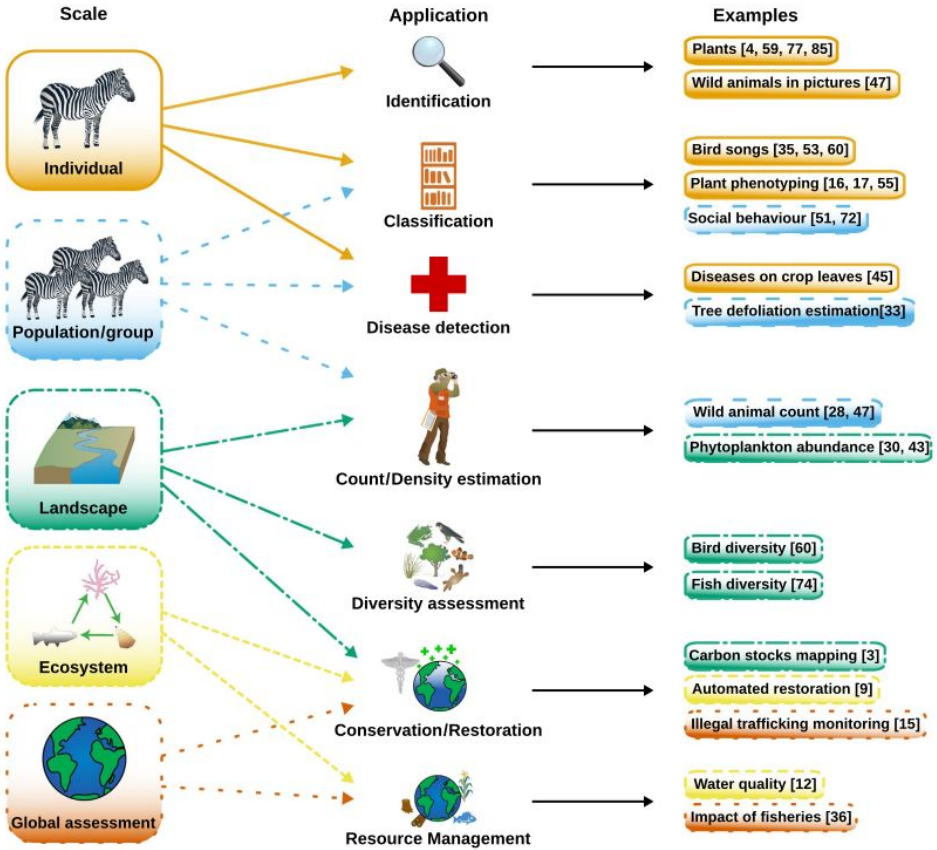

Christin et al., 2019

Pose tracking

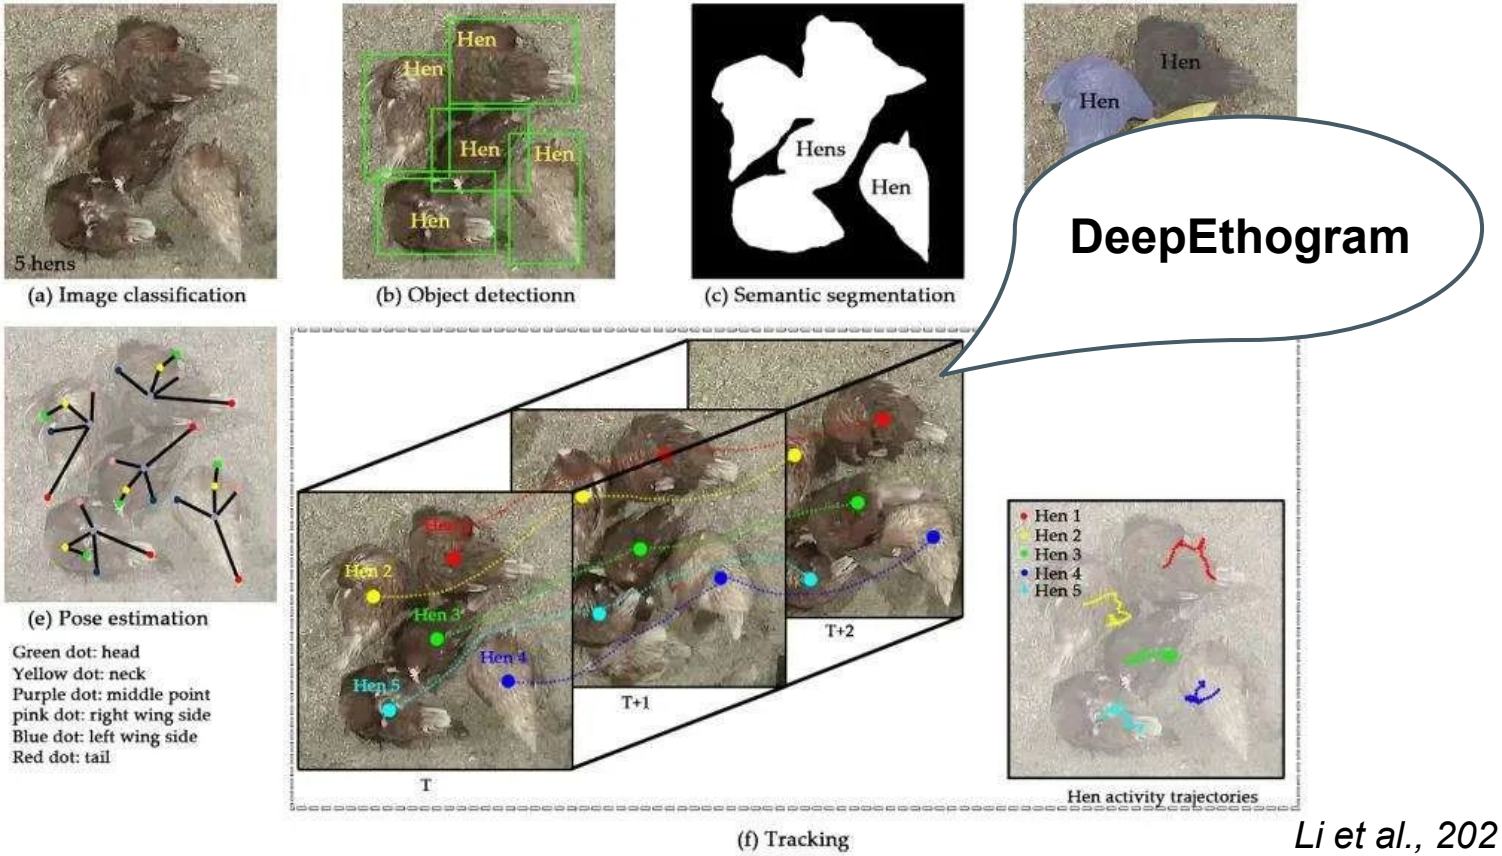

Li et al., 2021

Neural Style Transfer

NST = Process of using CNNs to render a content image in different styles

Neural Style Transfer

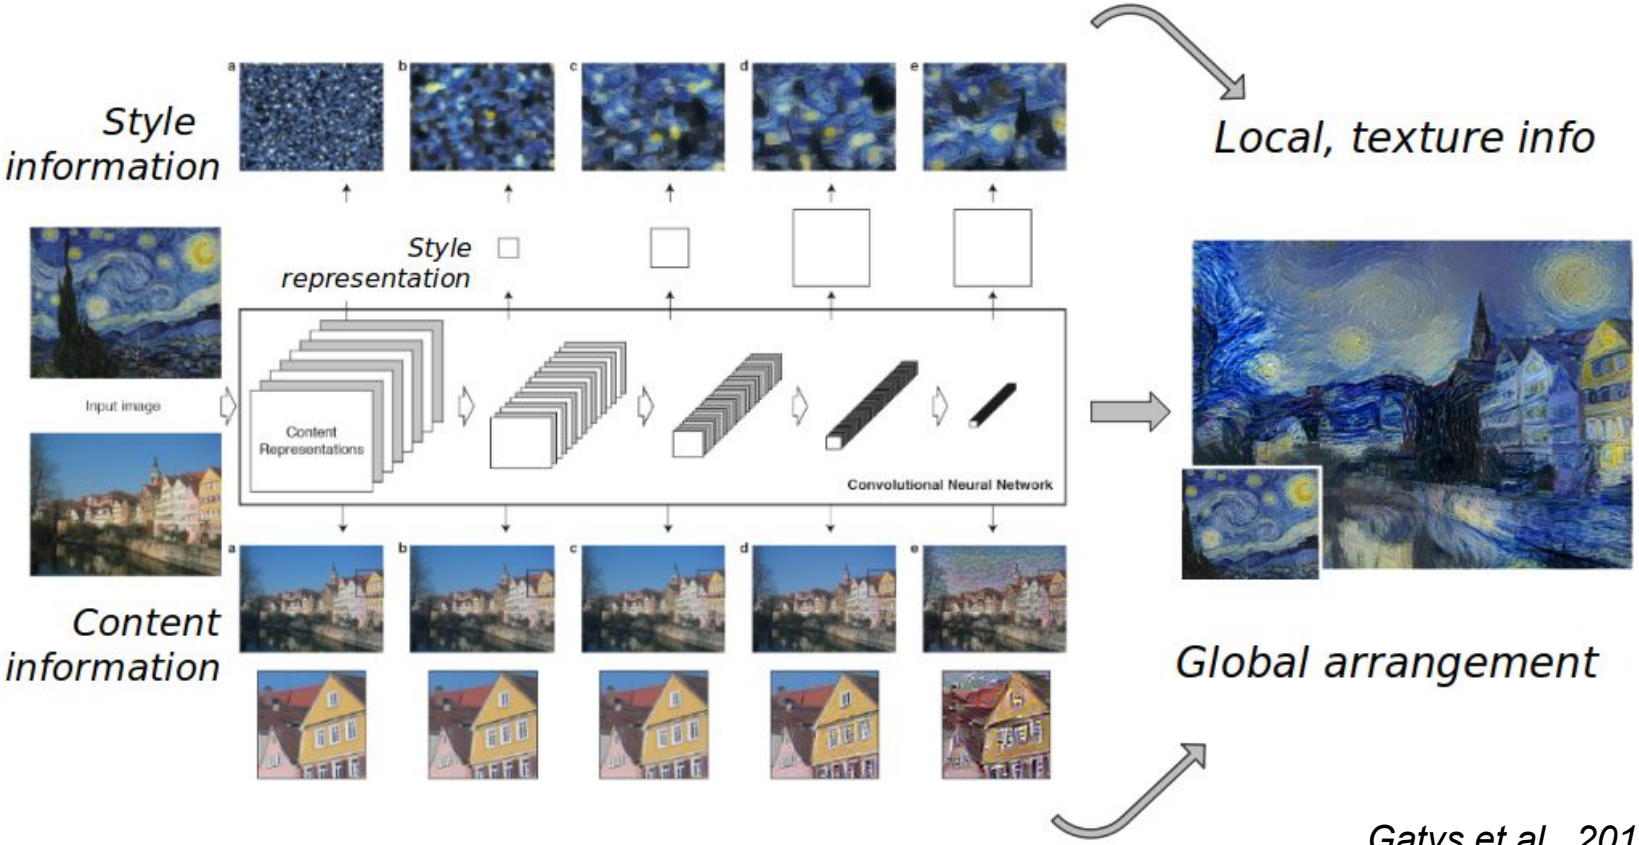

Gatys et al., 2016

Neural Style Transfer

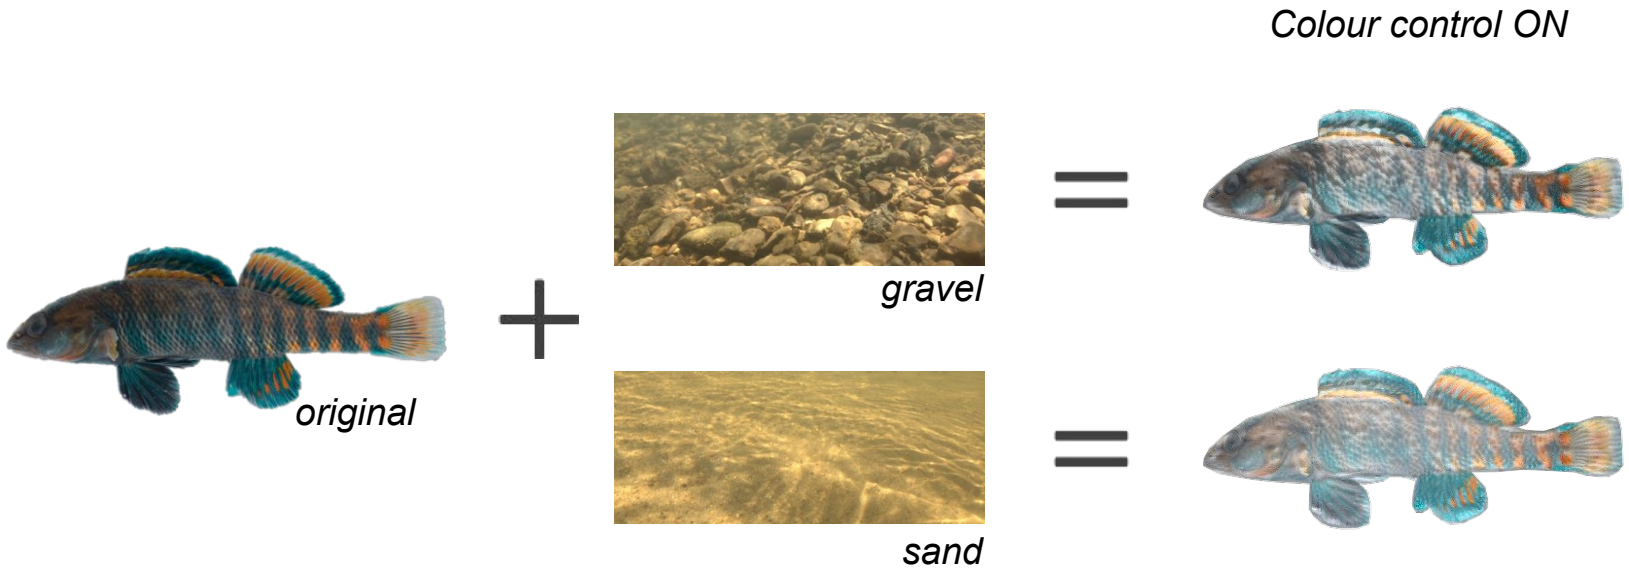

Deep learning: Further applications

- CamoGAN: exploiting GANs to simulate an evolutionary arms race between the camouflage of a synthetic prey and its predator *(Tálas et al., 2019)*
- ButterflyNet: using a CNN to provide a comprehensive quantification of visible phenotypic similarity and an objective test of taxonomic delimitation *(Cuthill et al., 2019)*
- VGG-Mandrill: estimating facial resemblance using a CNN and investigating its link with kin selection *(Charpentier et al., 2020)*
- **Exploiting the similarities between artificial and biological neural networks**