Deep Neural Networks

Joana Frontera-Pons

March 3rd, 2020 - Project THINK - Kick-off meeting

OUTLINE OF THE TALK

Objectives

- Propose and set the **goals** for the training,

Organisation

- Decide on the **organisation** of the different training sessions,
- Discuss the number, the length and modalities of the sessions.

Content

- Introduce the different **topics** that will be addressed during the training,

Discussion

- Open discussion, questions comments and remarks

OBJECTIVES OF THE TRAINING

- Learn the basic bricks of Artificial **Neural Networks** and Deep Learning,
- Overview of the main **optimisation** strategies in deep learning,
- Understand how to evaluate and improve the performance of a Deep Neural Network,
- Get familiar with **Convolutional Neural Networks** and learn how to use the most widely spread models,
- Use CNNs as building blocks for different types of architectures (ResNets, YOLO,…),
- Propose new architectures adapted to the applications in the THINK project.

ORGANISATION

Theoretical Sessions

- 10 theoretical sessions (1h30 each)
- Webminar through 5 or 10 weeks
- Alternatively, concentrated in one week ?
- Is the content appropriate?

Exercises

- Proposal of exercises/mini-project through the training
- Implementation of the different topics in Python

SESSION 1: INTRODUCTION

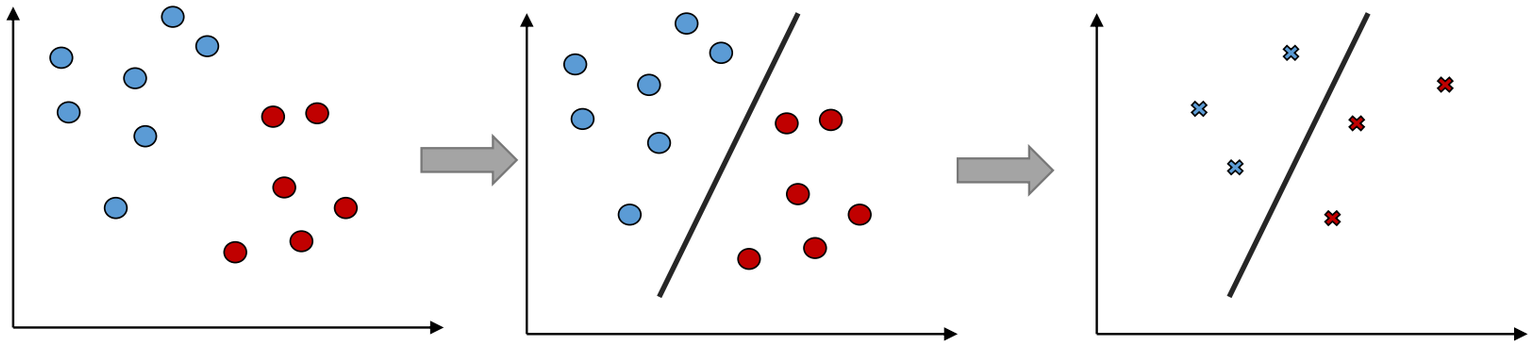

Supervised

Shallow	Deep
• Support Vector Machine • Logistic Regression	• Convolutional Neural Net • Recurrent Neural Net
• Denoising Autoencoder • Restricted Boltzmann machines • Sparse coding	• Stacked Denoising Autoencoder • Deep Boltzmann machines • Hierarchical Sparse Coding

Unsupervised

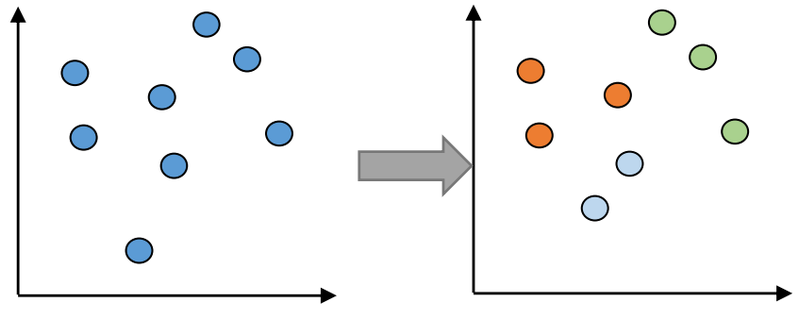

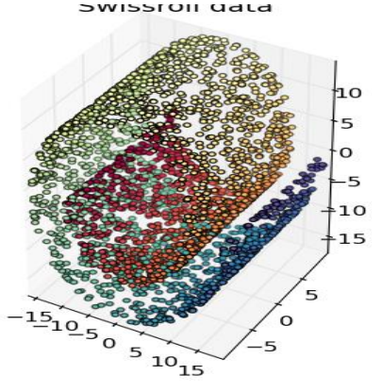

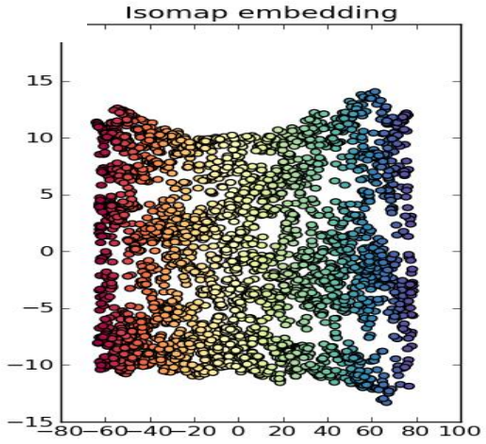

SESSION 2: DEEP NEURAL NETWORKS

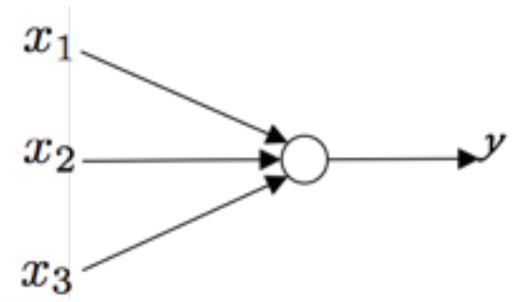

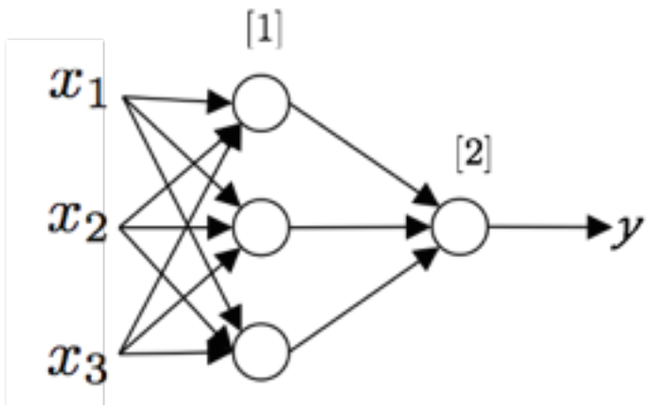

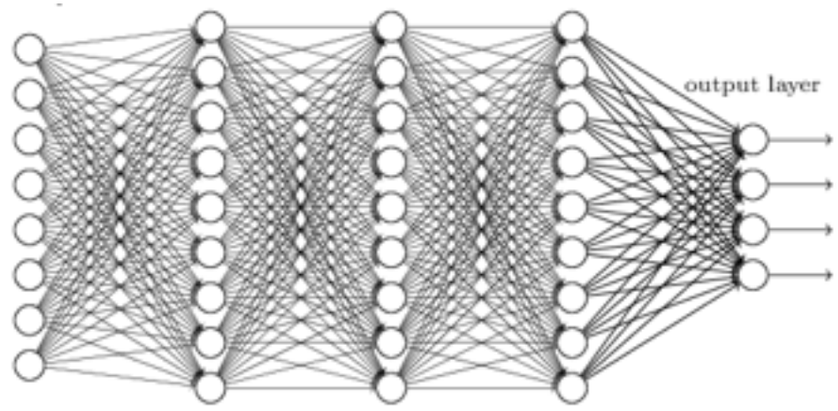

Shallow

Deep Neural Network

Need to define:

- L : number of layers,
- $n^{[l]}$: number of nodes (units) in layer l,
- $a^{[l]}$: *activations* in layer l,
- $a^{[L]}$: Output/ prediction

Session 3: Bias and Variance / Regularization

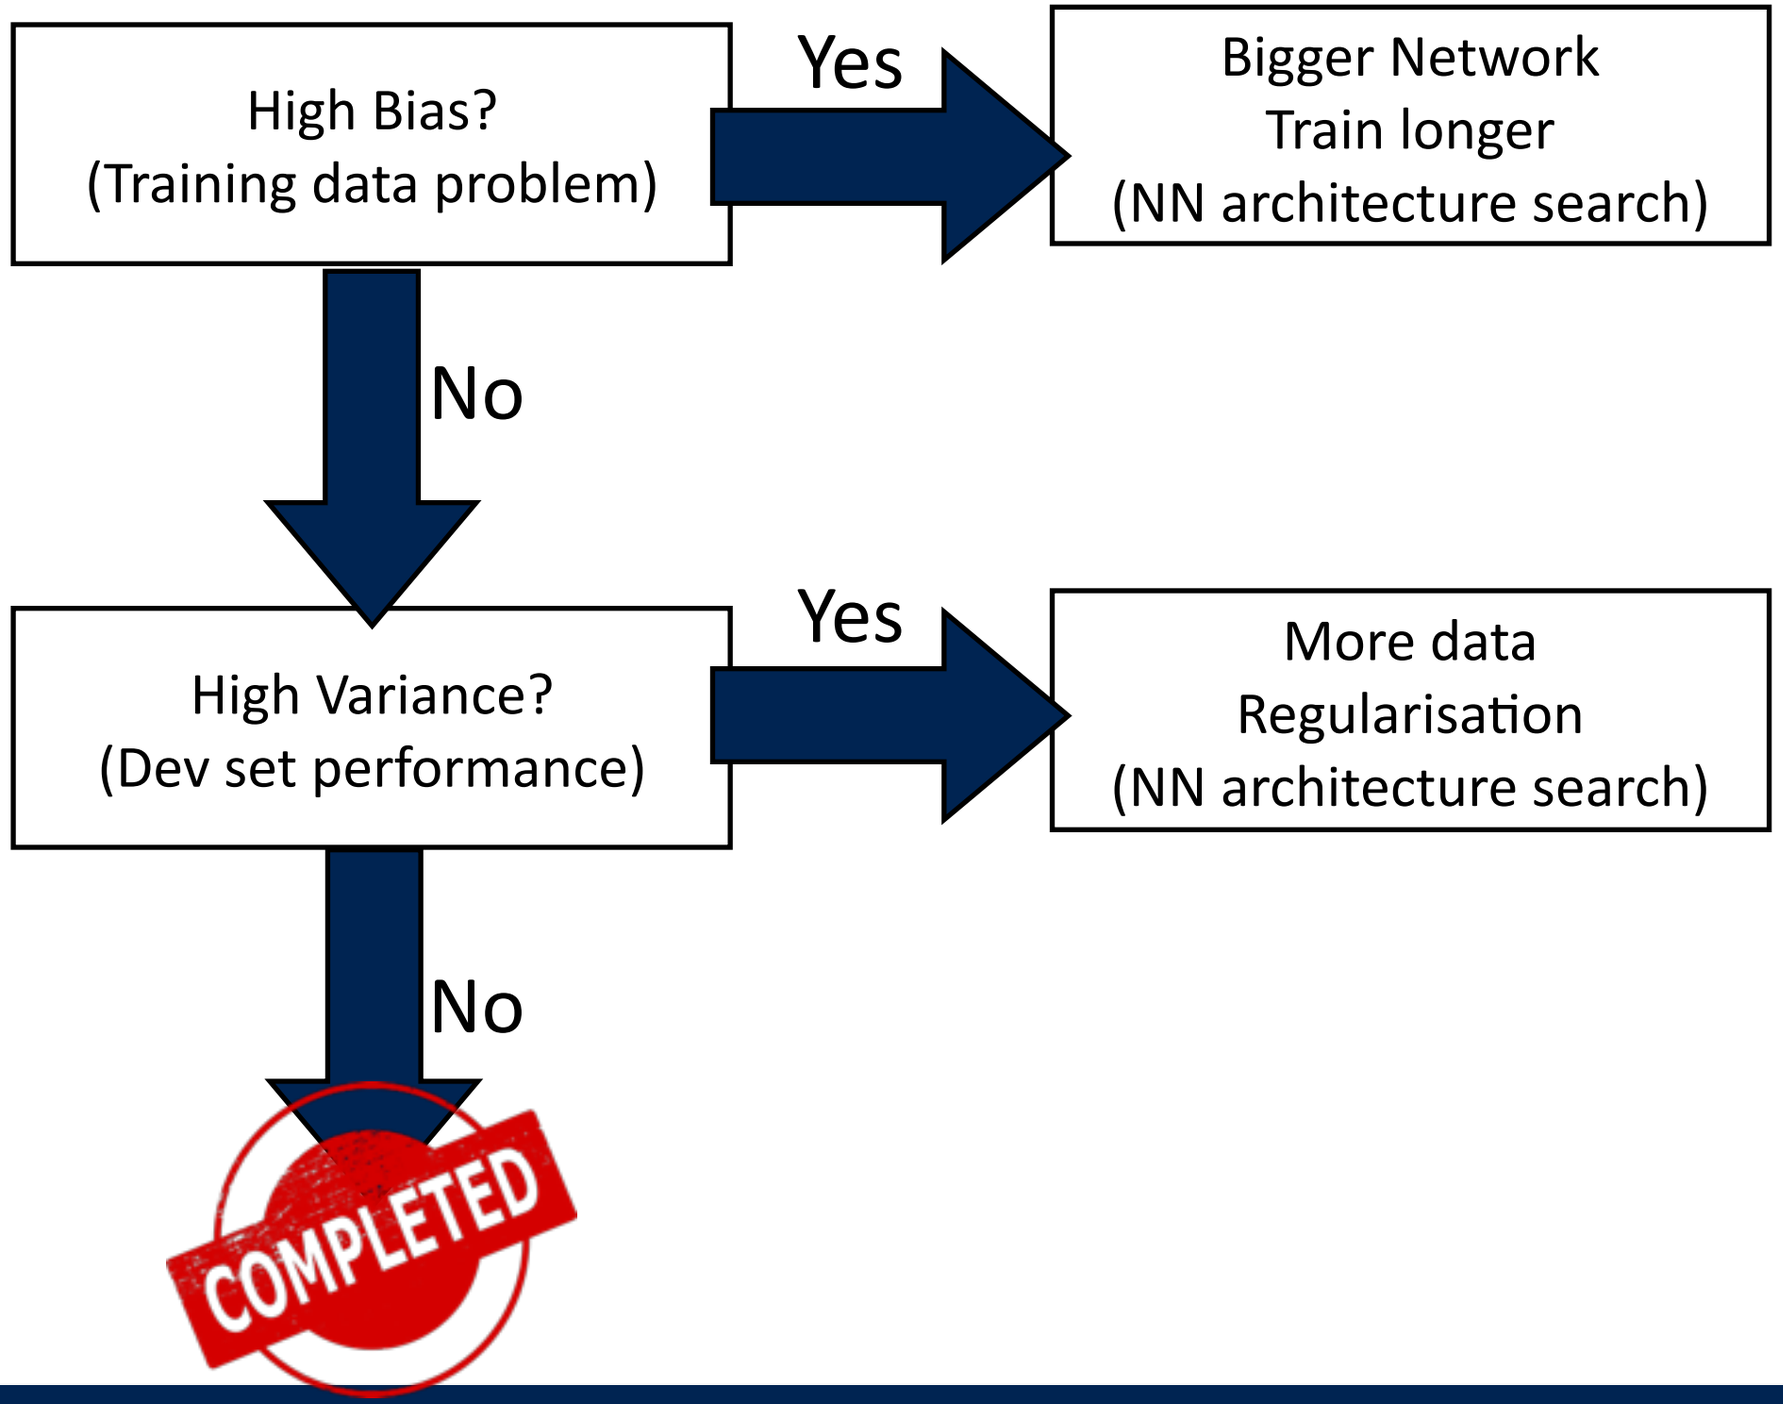

Session 4: Optimisation Strategies

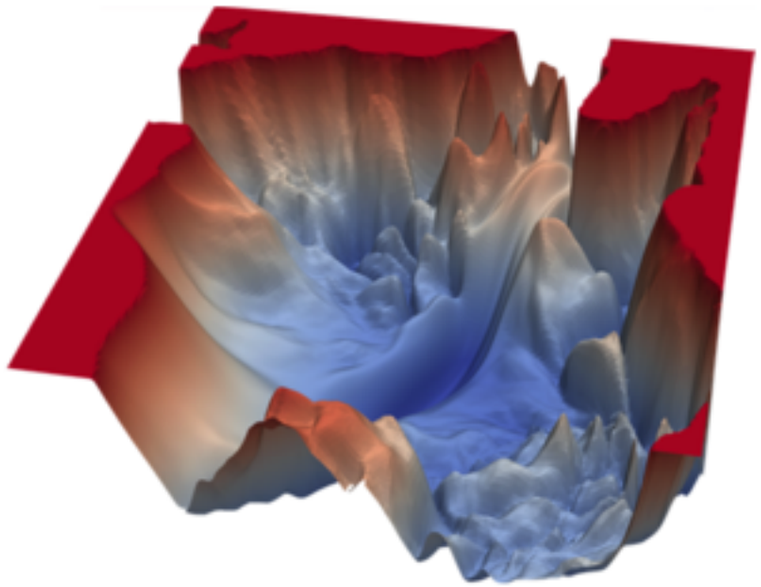

Figure: Loss Function for VGG56 [Li et al., 2017]

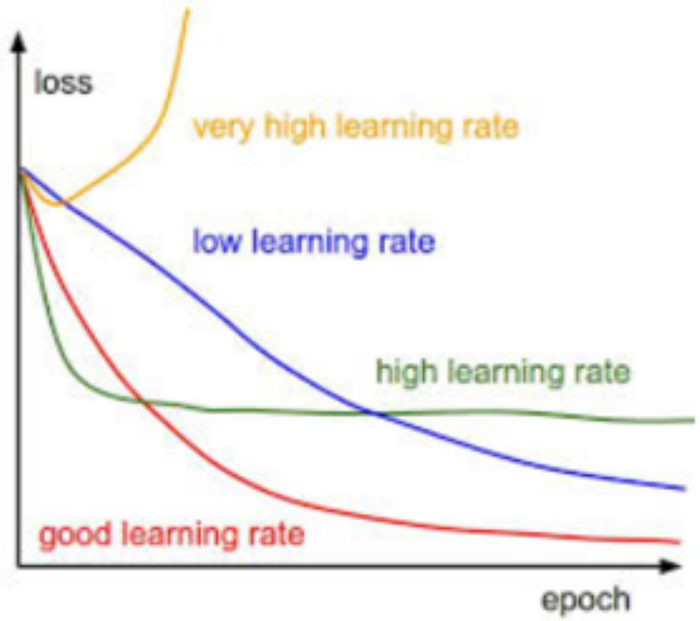

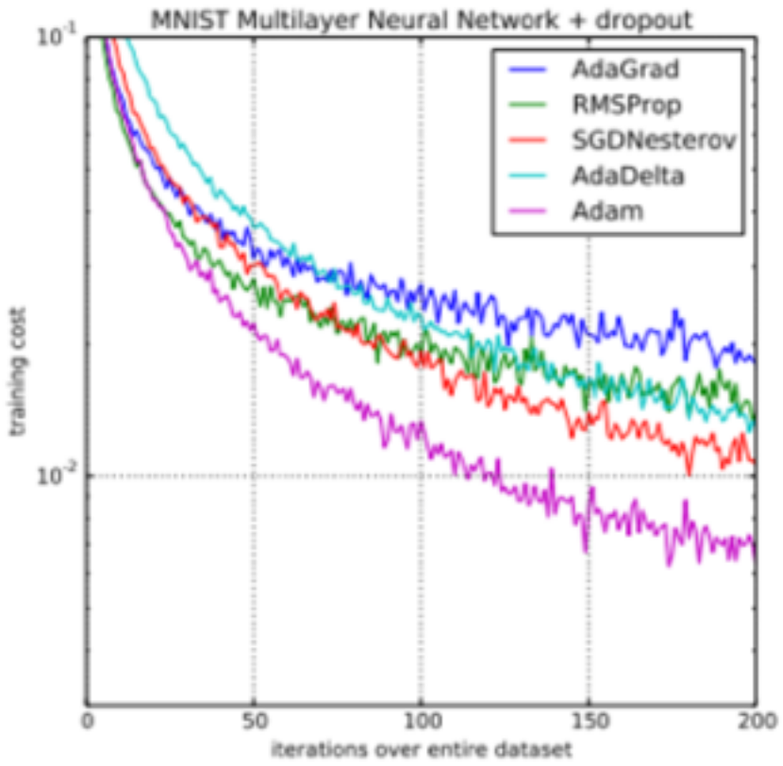

SESSION 5: RECURRENT NEURAL NETWORKS

Examples:

- Speech recognition,
- Sentiment analysis,
- DNA sequence analysis,
- Machine translation

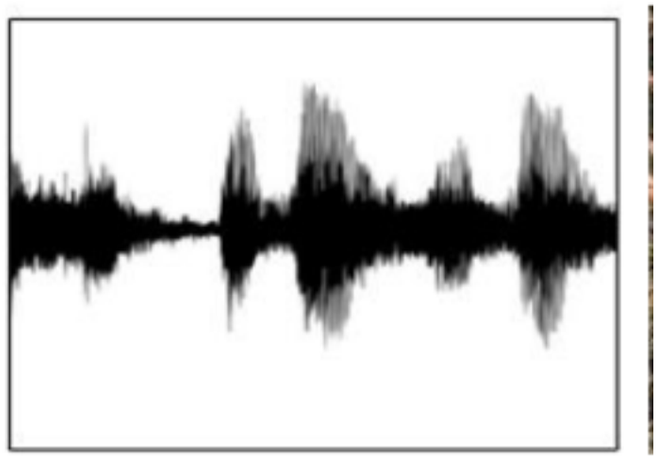

Audio

Four scores and seven years ago...

Text

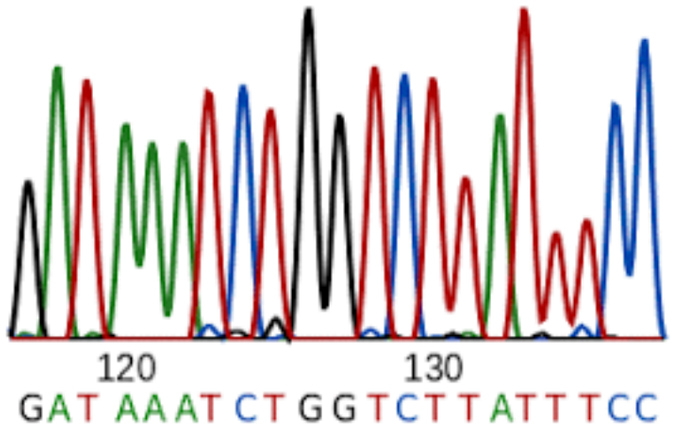

Session 6: Convolutional Neural Networks

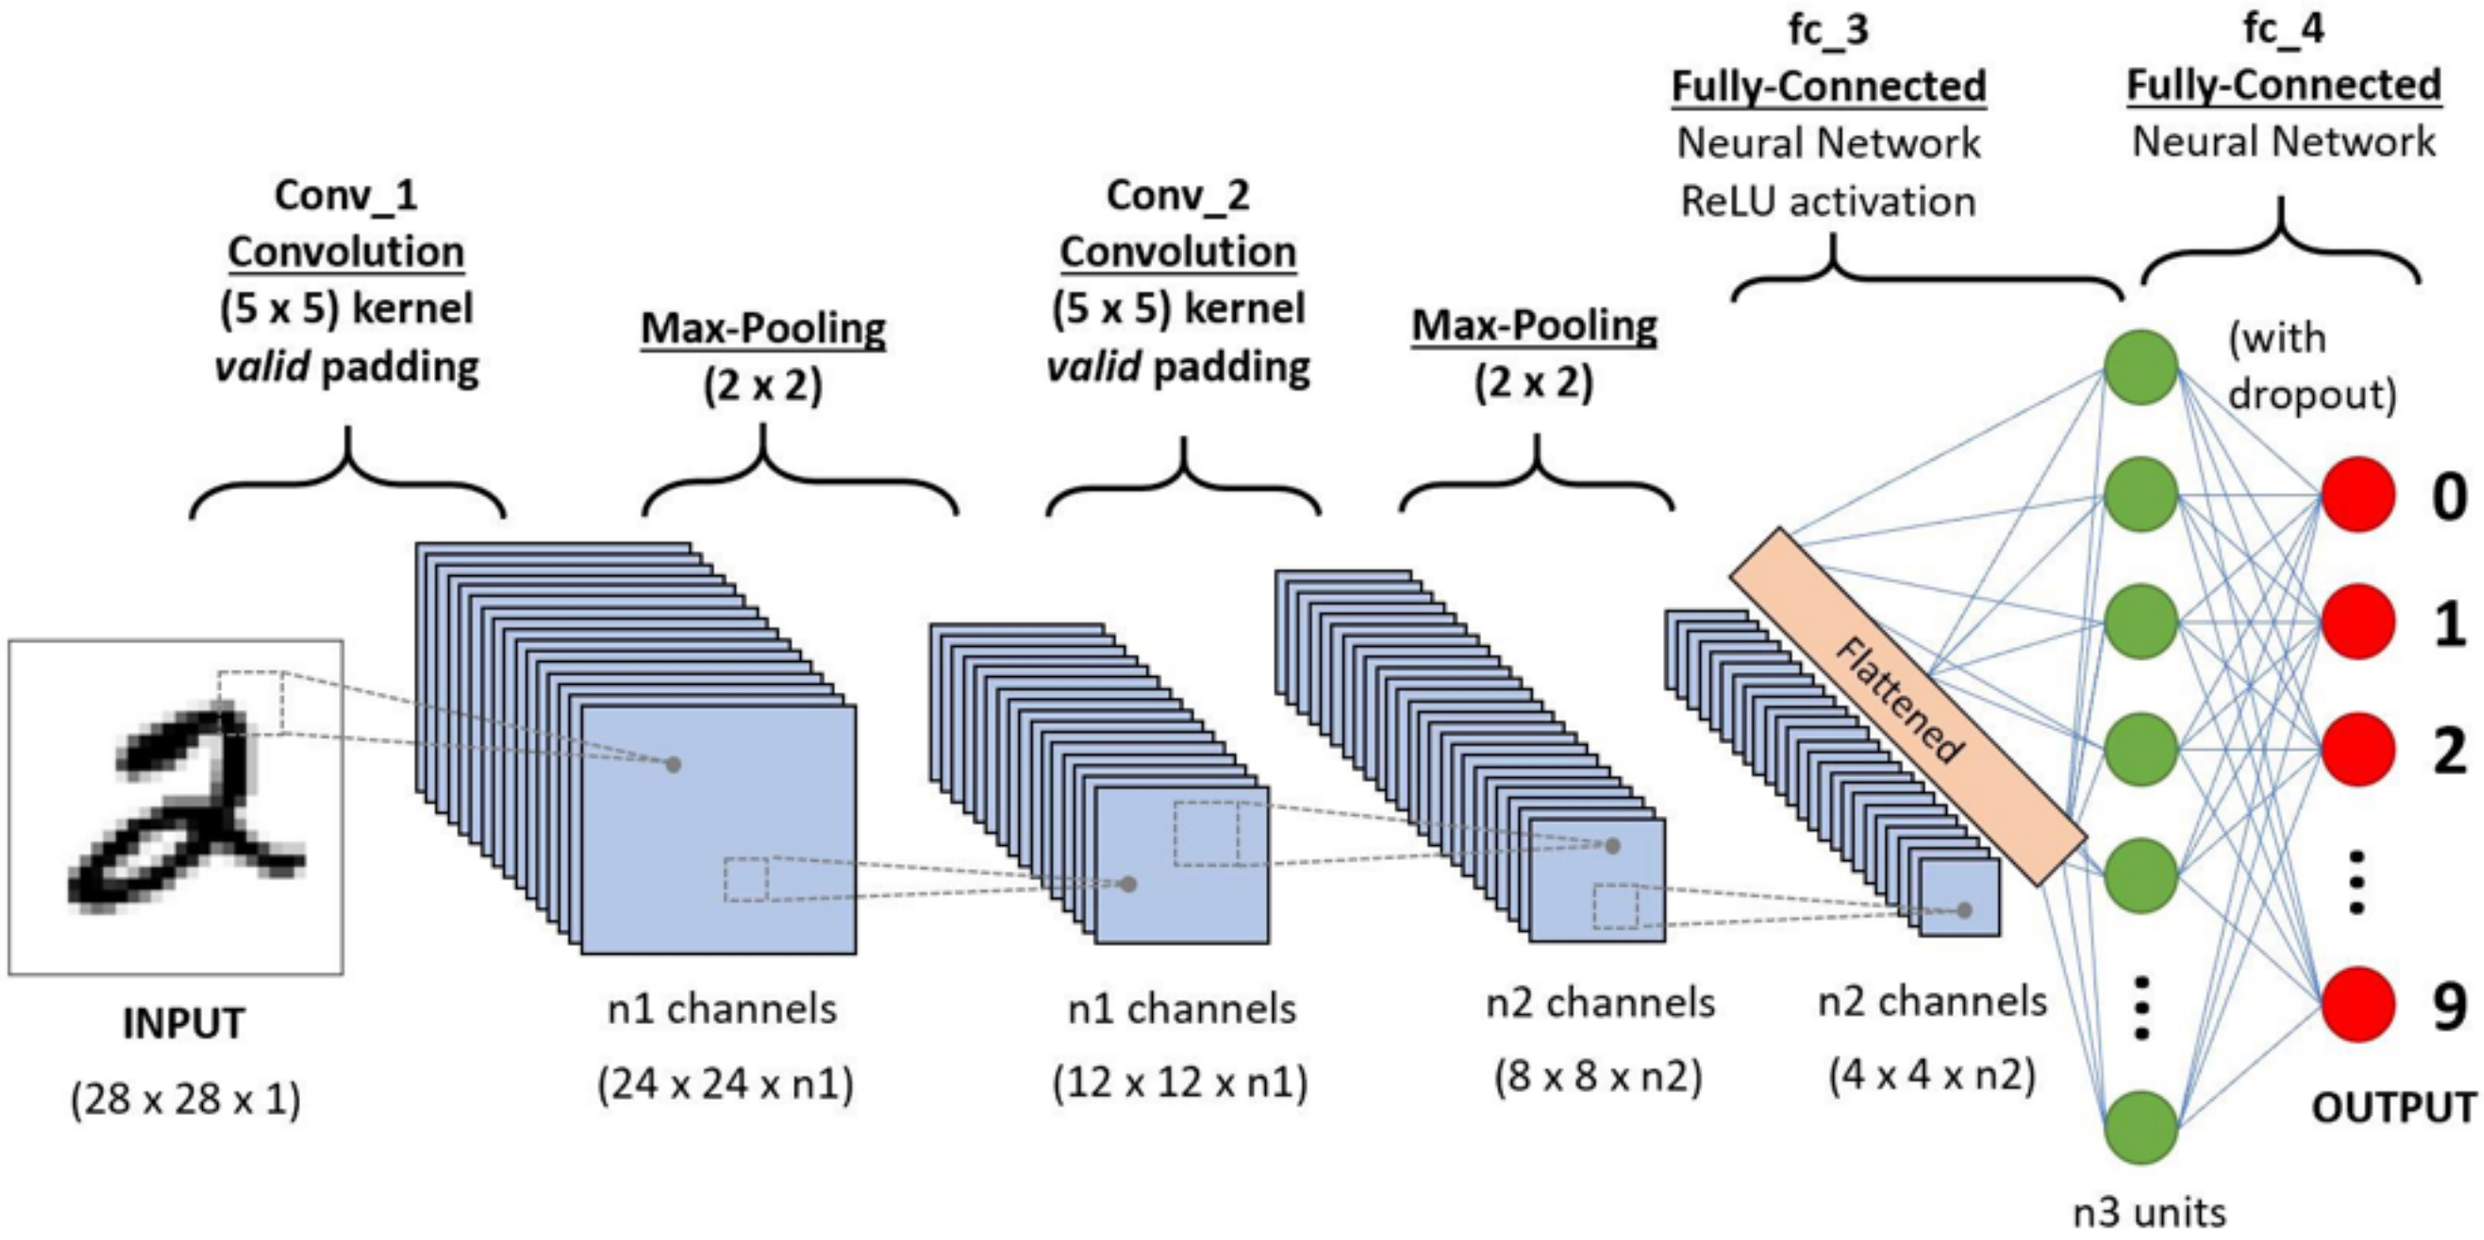

SESSION 7: COMMON CNN ARCHITECTURES

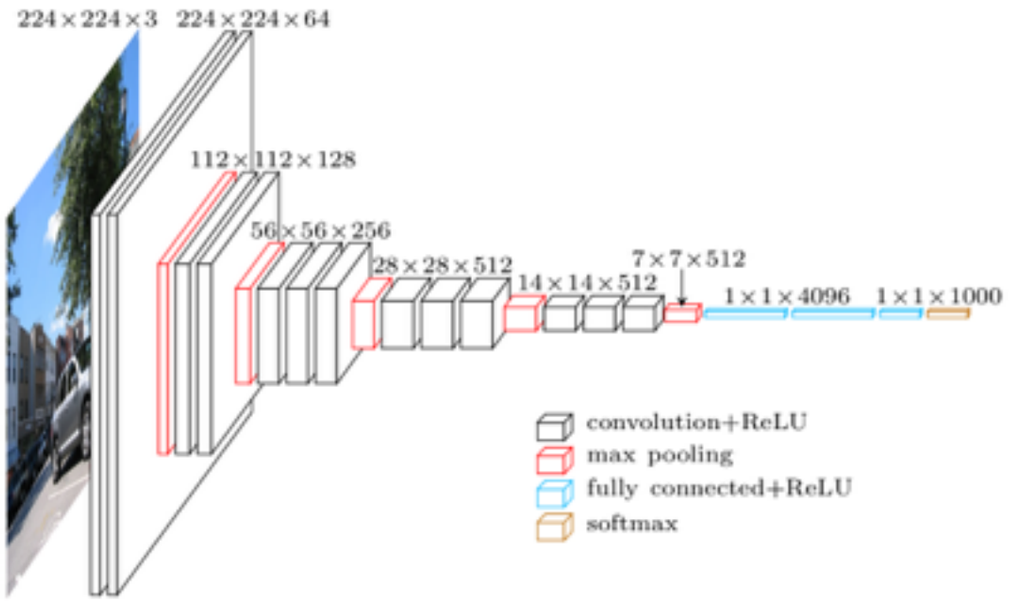

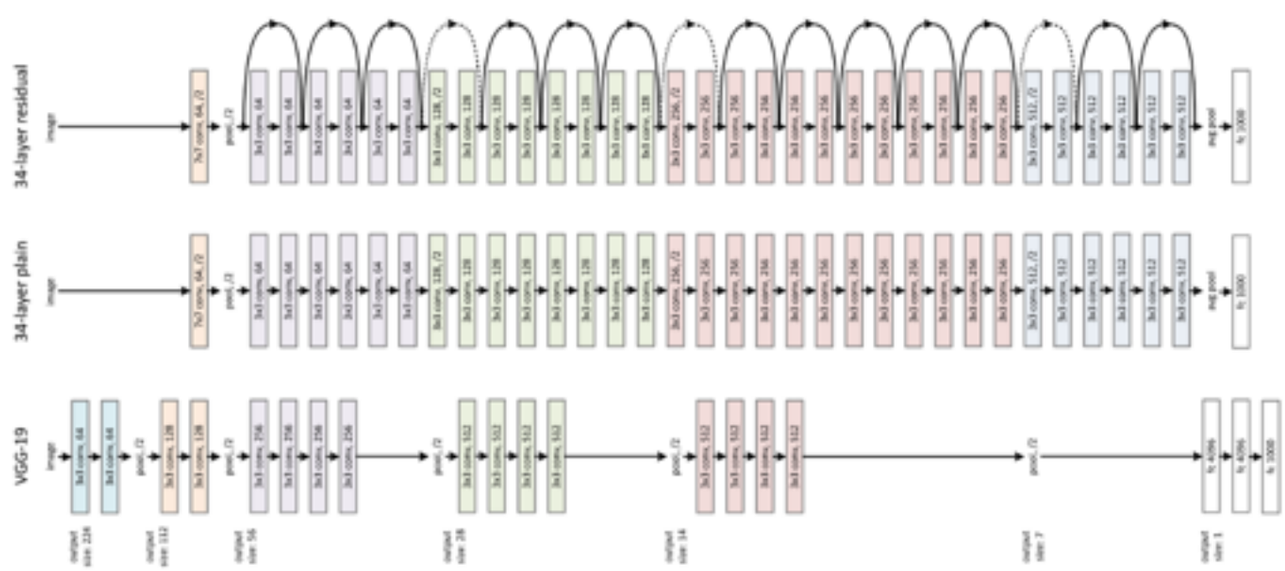

Simonyan, Karen, and Zisserman. "Very deep convolutional networks for large-scale image recognition." (2014)

VGG - 16

Residual Networks

Application to One-shot learning and Style Transfer

Session 8 : Object Detection and Segmentation

- Classification: There is a balloon in this image.
- SemanticSegmentation: These are all the balloon pixels.
- ObjectDetection:There are 7 balloons in this image at these locations. We're starting to account for objects that overlap.
- Instance Segmentation: There are 7 balloons at these locations, and these are the pixels that belong to each one.

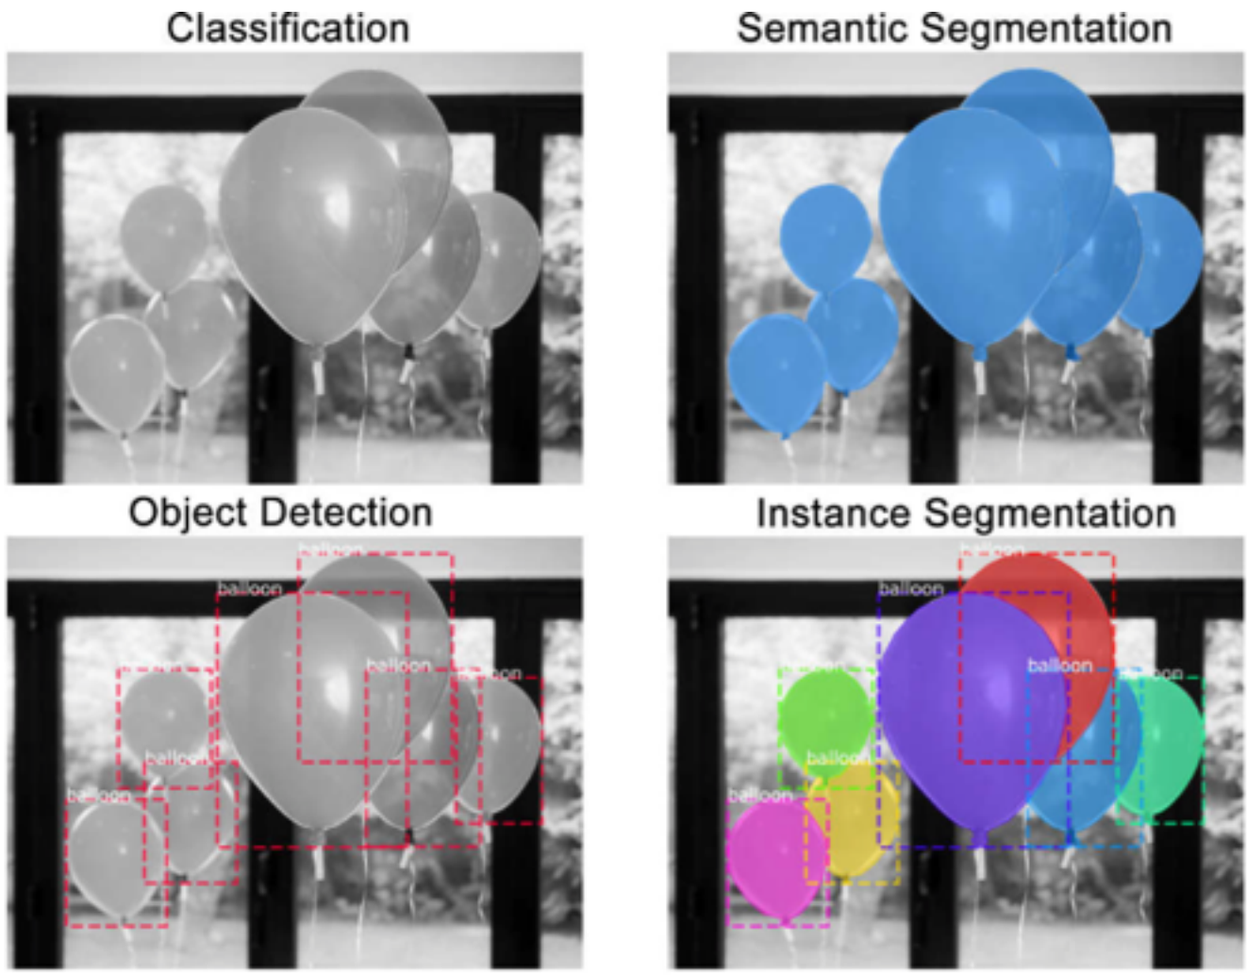

Instance segmentation is the task of identifying object outlines at the pixel level.

SESSION 9: REPRESENTATION LEARNING

- Train the model in order to **reconstruct** as accurately as possible the input.

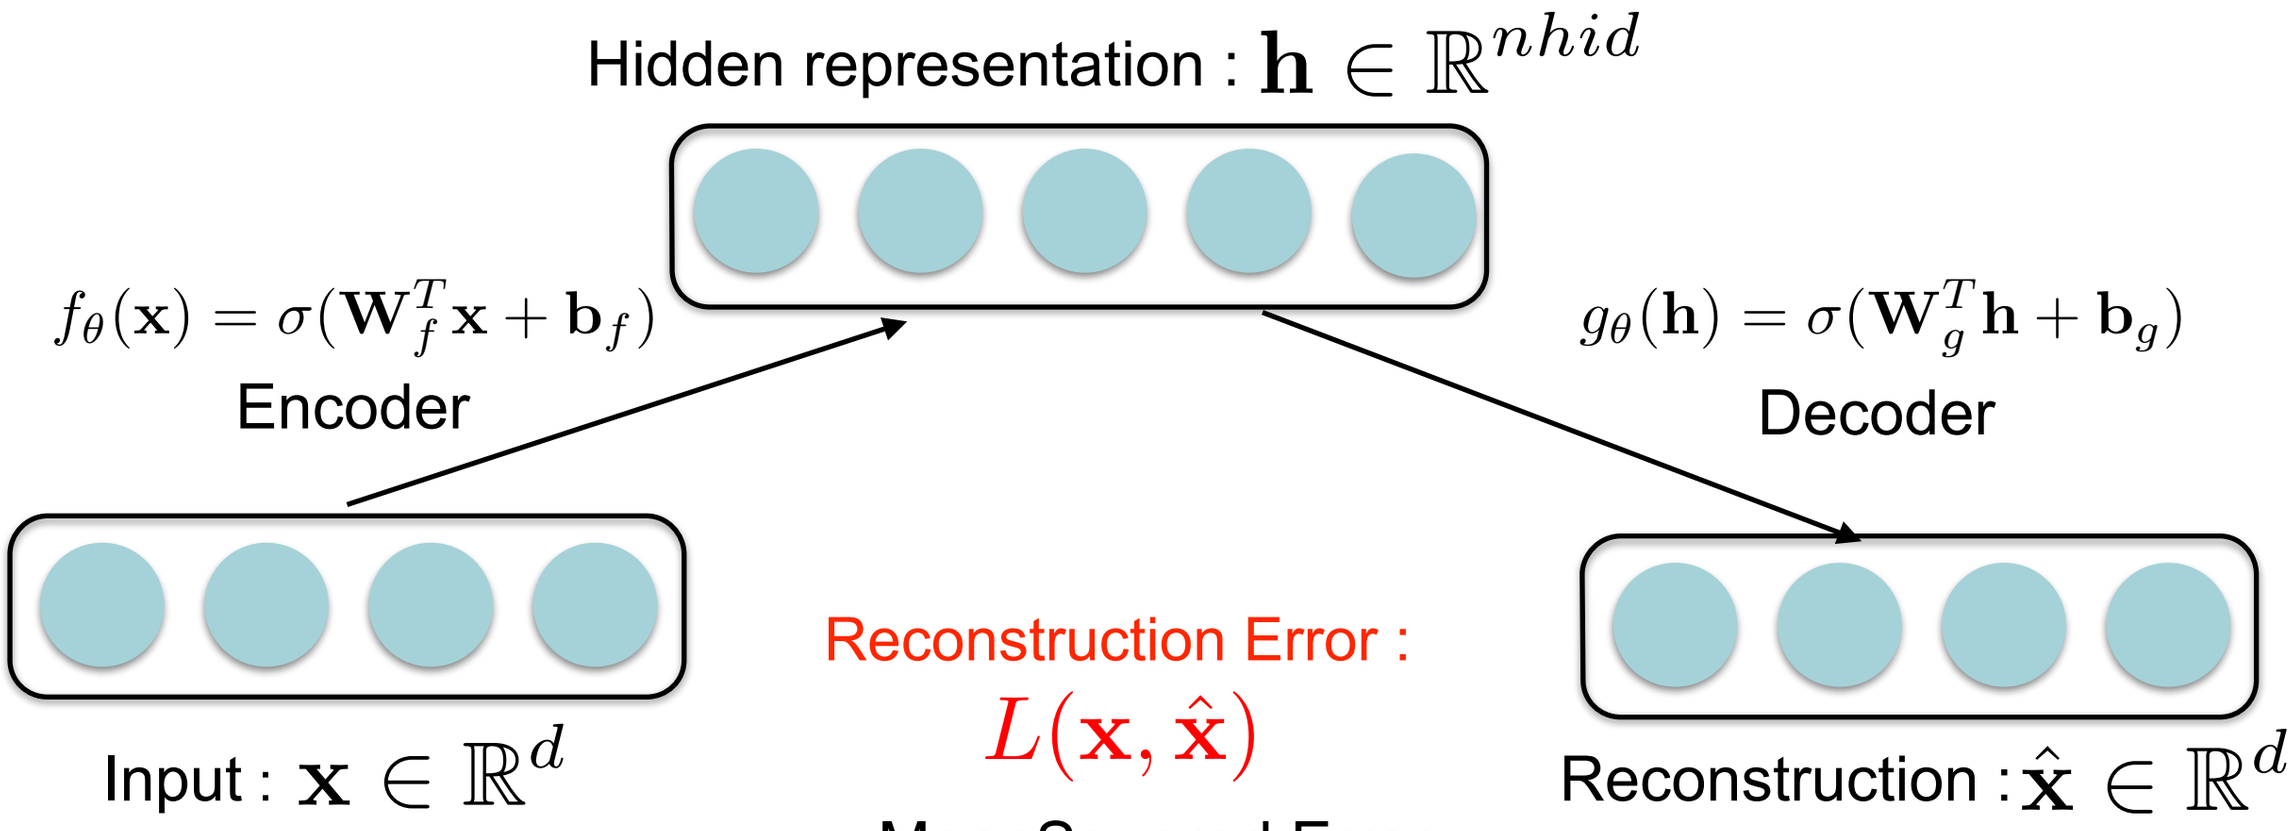

SESSION 10 : GENERATIVE MODELS

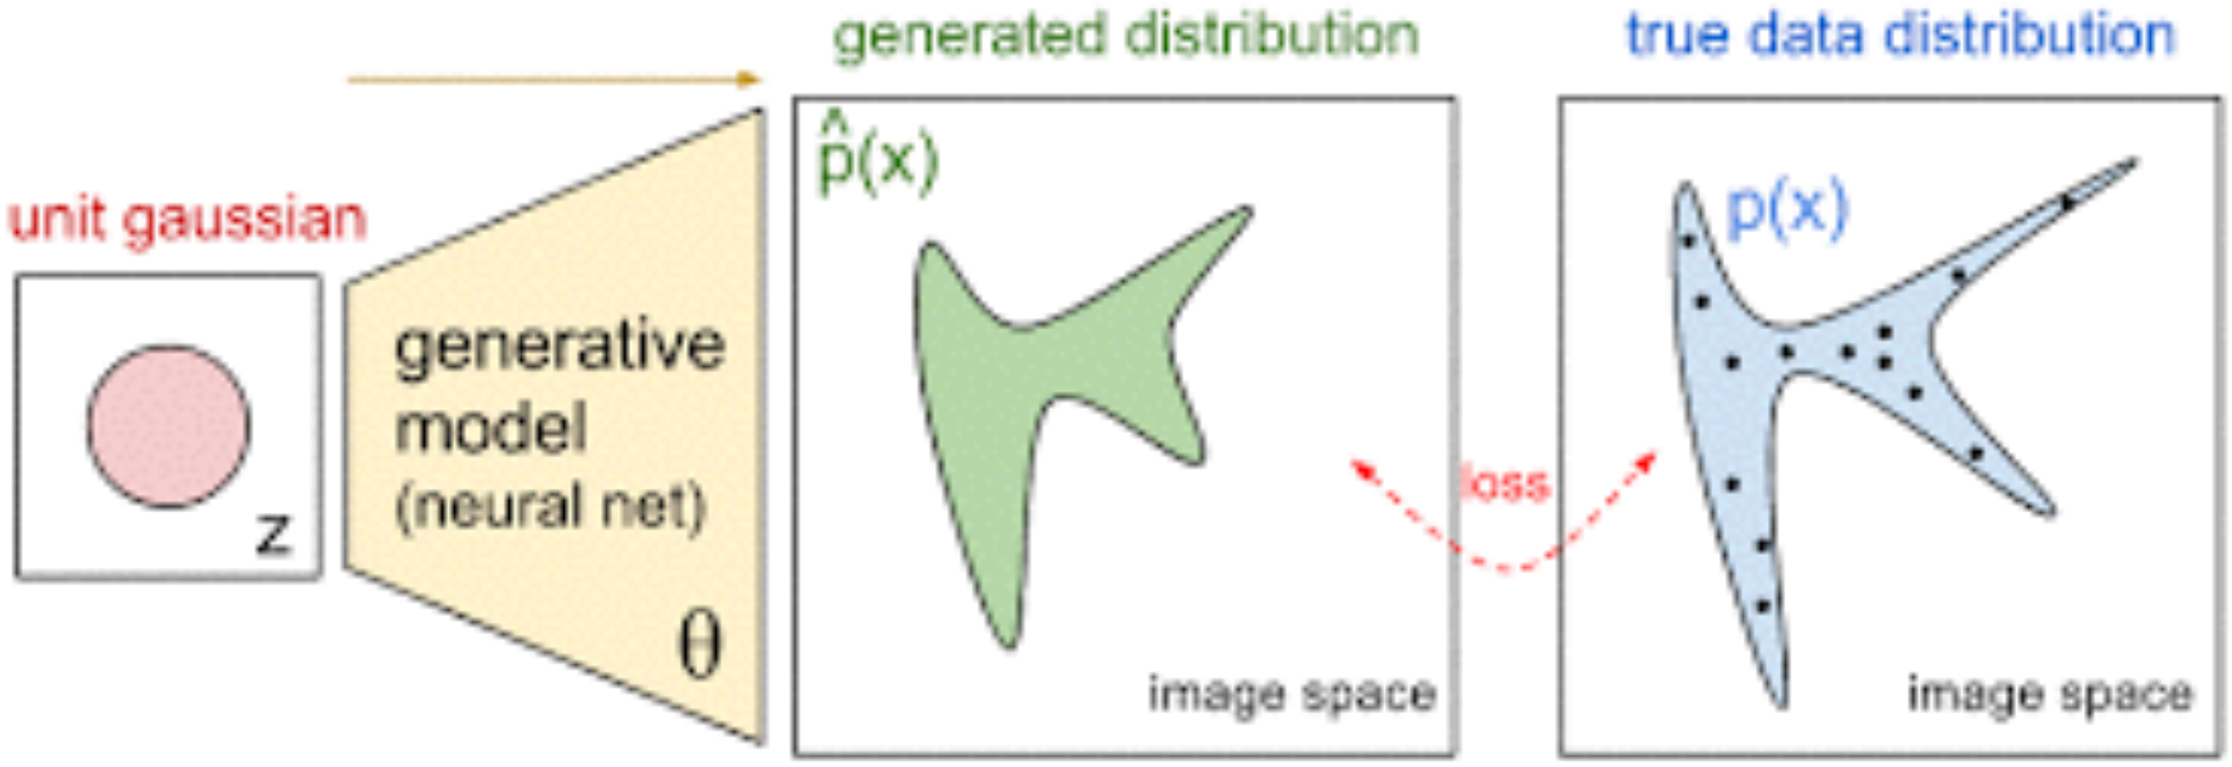

- **Advantatges and inconvenients of:**
  - **Variational auto-encoder**
  - **Generative Adversarial Networks**

Thank you for your attention!