Artificial Neural Networks and Deep Learning (Syllabus)

The following is a list of topics, questions, subjects which you are expected to know at the end of the course on Artificial Neural Networks and Deep Learning. For this reason you can use it to double check your preparation before the written exam.

Please recall the exam is meant to verify what you have understood from the lectures provided by the teachers, having said this, note:

- This is NOT the list of questions and exercises you will find in the exam, but we will look at this list when preparing it, so the exam questions are going to check these points
- If you know how to answer/face any of the following then you know how to face any possible question for the exam

Machine vs Deep Learning

- What is machine learning
- What are the machine learning paradigms and their characteristics
- What are the machine learning task and their characteristics
- What is deep learning and how it differs from machine learning
- Examples of successful deep learning applications
- What is transfer learning and what it is used for
- What is a feature and its role in machine and deep learning
- What is cross-validation and what it is used for
- What are hyperparameters and identify those for all the models presented in the course
- What is Maximum Likelihood estimation and how does it work in practice
- What is Maximum a Posteriori and how does it work in practice

Feed Forward Neural Networks

- The Perceptron and its math
- The Hebbian Learning paradigm, its maths, and it application
- The feed forward neural architecture and its improvement wrt the Perceptron
- The number of parameters of any model you can build
- Activation functions, their math, their characteristics, and their practical use
- What is backpropagation and how does it work
- The relationship with non-linear optimization and gradient descent
- Backpropagation computation for standard feedforward models
- Online vs batch, vs mini-batch learning and their characteristics
- Forward-backward algorithm for backpropagation
- Derivatives chain rules
- Error functions, their statistical rationale, their derivation, and their derivative
- The issue of overfitting and its relationship with model complexity
- Regularization techniques, their rationale and their implementation
- Techniques for hyperparameters tuning and model selections
- Techniques for weights initialization and their rationale
- Batch-Normalization rationale and use

Convolutional Neural Networks

- Convolution and correlation, linear classifier for images
- The layers of a CNN
- Connections between CNN and Feedforward-NN, interpretations of CNN
- How to compute the number of parameters of a CNN
- Best practices to train CNN models: data augmentation, Transfer learning
- Design criteria from the most successful architectures shown during classes (no need to know exactly the architectures)
- Fully convolutional CNN and CNN for segmentation
- The key principles of CNN used for object detection
- Residual learning
- GANs, what they are and how they are trained

Recurrent Neural Networks

- Models for sequence modeling
- The architecture and math of a recurrent neural network
- Backpropagation through time rationale and math
- The limits of backpropagation through time and the vanishing/exploding gradient issue
- The vanishing gradient math in a simple case
- The Long Short-Term Memory model, rationale, ad math
- The Gated Recurrent Unit
- LSTM Networks, hierarchical and bidirectional.

Sequence to Sequence Learning

- Sequential Data Problems, with examples
- The Seq2Seq model, training, and inference
- Neural Turing Machine model and the attention mechanism
- Attention mechanism in Seq2Seq models
- Chatbot: core models and context handling
- The Transformer model

Word Embedding

- Neural Autoencoder model, error function, and training
- Language models and the N-grams model
- Limits of the N-gram model
- The concept of embedding and its benefits
- The Neural Language Model and its use for word embedding
- Google’s word2vec model (CBOW)
- Possible uses of word embedding