Neural Networks: Theory and Applications (Prof. Dr. Stefanie Vogl)

1. Overview

The Neural Networks course provides a comprehensive introduction to the principles and applications of neural networks. It covers fundamental concepts, such as the architecture and training of neural networks, and delves into advanced topics including convolutional and recurrent neural networks. Students will learn to implement neural networks using popular frameworks like Keras and will explore various real-world applications, such as image recognition and time series forecasting. The course emphasizes hands-on experience through practical exercises and projects, enabling students to apply theoretical knowledge to practical problems. By the end of the course, students will be equipped with the skills needed to design, implement, and evaluate neural networks in diverse contexts.

2. Prerequisites:

- Basic programming knowledge (e.g., Python)
- Basic knowledge of linear algebra and probability theory
- Foundations of statistics and machine learning are advantageous

3. Learning Objectives:

Upon completion of the module, students should be able to:

- Understand the fundamental concepts and structures of neural networks.
- Know different types of neural networks and their areas of application.
- Implement neural networks using Python and popular libraries (e.g. Keras).
- Evaluate and optimize the performance of neural networks.
- Contextualize current research questions and developments in the field of neural networks.

4. Content:

- Introduction to Neural Networks

    - Historical overview
    - Biological inspiration
    - Basic terms and concepts

- Neural Network Architecture

    - Neuron model
    - Layer architectures (simple perceptron, multilayer perceptrons)
    - Activation functions

- Training Neural Networks

    - Forward and backward propagation
    - Loss functions

- Optimization algorithms (gradient descent, stochastic gradient descent, Adam)

- Regularization and Optimization

  - Overfitting and underfitting
  - Regularization techniques (L1, L2, dropout)
  - Hyperparameter tuning

- Special Network Architectures

  - Convolutional Neural Networks (CNNs)
  - Recurrent Neural Networks (RNNs)
  - Long Short-Term Memory Networks (LSTMs)
  - Auto-Encoders

- Applications of Neural Networks

  - Image recognition
  - Time series analysis

- Practical Implementation

  - Introduction to Keras (sequential models and functional API)
  - Implementing and training a simple neural network
  - Case studies and projects

5. Teaching and Learning Methods:

- Lectures with theoretical and practical content
- Practical computer exercises
- Project work and case studies
- Group work and presentations

6. Literature:

- "Deep Learning" by Ian Goodfellow, Yoshua Bengio, and Aaron Courville
- "Neural Networks and Deep Learning" by Michael Nielsen
- "Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow" by Aurélien Géron
- Scientific articles and current research papers (to be announced during the course)

7. Notes:

The module places a strong emphasis on practical applications and the implementation of neural networks. Students are encouraged to bring their own projects and questions from their areas of interest. Regular exercises and projects are essential to deepen theoretical content and develop practical skills.