MLNAT Protokolle

August 2022:

- Given a picture of 2D space with data points. What split achieves perfect purity?
- Calculate Bayes Rule: Population of 99% farmers and 1% librarians. 5% of farmers enjoy Sci-Fi and 90% of librarians enjoy Sci-Fi. Given someone enjoys Sci-Fi what are approximate probabilities that they are a farmer or a librarian.
- Which plot shows the ReLU function?
- What metric is good for testing for a rare but very deadly disease → Sensitivity
- What is the maximum entropy and when is it reached
- Advantages of transfer learning
- Advantages of CNNs
- How many learnable weights does a polling layer have
- By how much is an image decreased after going through a convolutional layer with some parameters and through pooling with some other parameters afterwards.
- Design an ML pipeline for a scenario: Predicting band gaps from SMILES with some labeled data and a simulation tool
- What is the Markov property (but there are two answers correct, which is unexpected)
- Draw example SMILES code and explain how it works in general
- How do molecule fingerprints work? Are they useful as output of a generative model?
- Some multiple choice question about the target networks in deep q learning
- What are the two limitations of deep q learning and what can we do about that
- Do we need a readout function for a GNN when used for learning potentials of single atoms in a molecule
- How to represent molecules in a GNN (some weird question with hand wavy terminology)
- What does the radius in a fingerprint correspond to in a GNN
- What is attention and why is it useful for translation and chemical reaction prediction
- Explain the Bayesian learning algorithm and two applications in ML and the natural sciences for it
- Explain a possible idea for parallel Bayesian learning
- Query by committee: What is it and how is it used in enhancing the test set