Universiteit Leiden
The Netherlands

Learning from small samples

Kocaman, V.

Chapter 2

Preliminaries

2.1 Definition of Terms

This chapter serves as a brief descriptions of the terms mentioned within this dissertation. We should note here that this chapter does by no means offer an exhaustive overview of the aforementioned fields, as this would lie out of the scope of the present thesis. Its aim, instead, is to acclimatize the reader to the definitions of terminology that will be recurring in the following chapters.

Machine learning (ML) Machine learning is a subfield of artificial intelligence that focuses on the development of algorithms and statistical models that can learn patterns and relationships in data. These algorithms can then be used to make predictions or decisions about new, unseen data.

Deep learning (DL) Deep learning is a subfield of machine learning that uses neural networks with multiple layers to model and solve complex problems. These deep networks are trained using large amounts of data and can learn hierarchical representations of data, allowing them to perform well on a variety of tasks such as image recognition, natural language processing, and speech recognition.

Overfitting Overfitting occurs when a machine learning model is too complex for the amount of data it is trained on. As a result, the model fits the training

data too closely and does not generalize well to new, unseen data, leading to poor performance.

Underfitting Underfitting occurs when a machine learning model is too simple and cannot effectively capture the underlying patterns in the data. As a result, the model will have poor performance on both the training and test data.

Self-supervised learning (SSL) Self-supervised learning is a form of unsupervised learning where the model learns from input data without explicit supervision, but with the goal of reconstructing the input data in some way. This can be done, for example, by predicting missing parts of an input or by rearranging the input data in a specific order.

Semi-supervised (SeSL) Semi-supervised learning is a form of machine learning that combines both supervised and unsupervised learning. In this setting, a small portion of the data is labeled, and the rest is left unlabeled which is usually large. The model then uses the labeled data for supervised learning and the unlabeled data for unsupervised learning to drive structure from unlabeled data to improve the overall performance of the task.

Unsupervised learning Unsupervised learning is a form of machine learning where the model learns from input data without explicit supervision. The goal is to find patterns or relationships in the data without being told what the target outputs should be. It relies on the model learning the structure of the input data based on the features present in the data via feature extraction, which involves identifying and extracting relevant characteristics from the raw data.

Softmax The softmax function is a mathematical function commonly used in machine learning as an activation function for multi-class classification problems. The function maps input values to a probability distribution over multiple classes, allowing the model to make predictions about the class with the highest likelihood.

Probably Approximately Correct (PAC) The Probably Approximately Correct (PAC) framework is a theoretical framework for understanding the generalization ability of machine learning algorithms. The framework provides a mathematical definition of when a learning algorithm is considered to have good generalization performance, making it a useful tool for comparing and evaluating different algorithms.

Synthetic data Synthetic data is artificially generated data used for training machine learning models. The data can be generated by a variety of methods, including simulation, extrapolation, and statistical sampling. The use of synthetic data can be useful in cases where obtaining real-world data is difficult or impossible.

Data Augmentation Data augmentation is a technique used to artificially increase the size of a dataset by generating new data from existing data. This can be done, for example, by applying random transformations or perturbations to the data, or by generating new samples from a generative model. The goal of data augmentation is to reduce overfitting by providing the model with more diverse training data.

Weight Decay Weight decay is a regularization technique in machine learning that penalizes large weights in the model to reduce overfitting and improve the generalization performance. It involves decreasing the magnitude of the weights over time, effectively reducing the complexity of the model and preventing it from memorizing the training data.

Dropout Dropout is a regularization technique used in deep learning to prevent overfitting by randomly dropping out neurons during training. This random dropout creates a different network architecture for each training iteration, forcing the network to learn more robust features.

Early Stopping Early stopping is a method in deep learning used to prevent overfitting by monitoring the performance of the model on a validation set, and stopping training once the performance on the validation set starts to degrade. This helps to ensure that the model is not overfitting to the training data.

Lottery Ticket Hypothesis The lottery ticket hypothesis is a concept in deep learning that suggests that sparse, randomly initialized neural networks can be trained to perform as well or better than dense networks. The hypothesis states that a small, well-initialized subnetwork of a larger network can be trained to perform well if it has the right connections and weights.

Ensemble Models Ensemble models are machine learning models that combine the predictions of multiple individual models to improve overall performance. The idea behind ensemble models is that by combining the predictions of several models, the strengths of each model can be leveraged to create a more robust overall system.

Transfer Learning Transfer learning is the process of using a pre-trained model on one task as a starting point for training on a different, but related, task. The idea is that the pre-trained model has already learned relevant features for the new task, allowing for faster and more effective training.

Regularization Regularization is a technique used in machine learning to prevent overfitting by adding a penalty term to the loss function during training. The goal is to encourage the model to learn a simpler and more general representation of the data.

Generalization Performance Generalization performance refers to a machine learning model's ability to make accurate predictions on new, unseen data. It is the ability of the model to generalize from the training data to unseen data, and is an important consideration in the design of machine learning algorithms.

Pruning Pruning is a method used to reduce the size and complexity of a machine learning model. The process involves removing unimportant connections or neurons from the model to make it more efficient and improve its generalization performance.

Bayesian Methods Bayesian methods are a family of machine learning algorithms that use Bayesian statistics to model the underlying relationships between inputs and outputs. Bayesian methods allow for the incorporation of prior knowledge and the estimation of uncertainty in the model.

Frequentist Methods Frequentist methods are a family of machine learning algorithms that use frequentist statistics to model the underlying relationships between inputs and outputs. It makes predictions on the underlying truths of the experiment, using only data from the current experiment (the parameters are fixed but unknown). Frequentist methods assume the observed data is sampled from some distribution. For example, in logistic regression the data is assumed to be sampled from Bernoulli distribution, and in linear regression the data is assumed to be sampled from Gaussian distribution.

Artificial General Intelligence (AGI) Artificial General Intelligence (AGI) is a field of artificial intelligence focused on the development of machine systems that have general cognitive abilities, including the ability to understand or learn any intellectual task that a human being can. It is a form of artificial intelligence that goes beyond narrow AI and aims to develop machines that can perform a wide

range of tasks.

L1 Norm Regularization (Lasso) L1 norm regularization, also known as Lasso, is a type of regularization in machine learning that adds a penalty term to the loss function that is proportional to the absolute value of the coefficients of the model. This regularization technique encourages the model to have sparse solutions, meaning that some of the coefficients will be exactly zero, effectively reducing the number of features the model uses. This can improve the interpretability and stability of the model, but may also reduce its accuracy.

L2 Norm Regularization (Ridge) L2 norm regularization, also known as Ridge, is a type of regularization in machine learning that adds a penalty term to the loss function that is proportional to the square of the magnitude of the coefficients of the model. This regularization technique encourages the model to have small, but non-zero coefficients, and helps to prevent overfitting by reducing the magnitude of the coefficients. The regularization term can be controlled by a hyperparameter, which determines the strength of the penalty, and the optimal value of this hyperparameter must be determined through cross-validation.