Fachpraktikum Interaktive Systeme: Machine Learning for Intelligent Mobile User Interfaces using TensorFlow

Winter term 2017/2018

Lecturer: Niels Henze / Huy Viet Le / Sven Mayer

Introduction

We currently witness the third wave of machine learning. In contrast to the previous waves, current machine learning research demonstrated impressive performance for a very wide range of tasks. Recent success in machine learning was mainly enabled by combining new training algorithms, new network architectures and moving the training on graphic cards. Combined, these three aspects enable to train models that not only outperform previous approaches but also enabled new application areas. Machine learning models are on par with humans or even demonstrate super-human performance for a number of tasks, including playing Go, playing Atari games, classifying images, or to determine the location where an image was taken. The human-computer interaction (HCI) community used machine learning for a very wide range of use cases. Amongst others, this includes classification of pro-eating disorder, authentication, touchscreen latency reduction, lifelogging, workout trainer, and chatbots. Training and using advanced machine learning models recently became much easier due to a variety of open libraries, including Torch, Theano, and TensorFlow. These libraries are not only used and developed by researchers from academia and industry but are also accessible for practitioners.

With increasing processing power, it became possible to train increasingly complex machine learning models. Luckily processing power is mainly needed during training. TensorFlow is a graph-based open source library for a wide range of machine learning algorithms. A unique feature of TensorFlow is the possibility to reduce the size of a trained model and compile it for deployment. In particular, it is possible to run models efficiently on mobile and embedded devices. This makes TensorFlow especially exciting for mobile HCI and Ubicomp researchers as it enables to use powerful machine learning models directly on end users devices. In this course, students will learn the basics steps of machine learning, develop neural networks and train them using TensorFlow. Further, students will learn how to port the trained model to a mobile phone using the model size reduction features of TensorFlow. They will apply the gained knowledge in practical group-projects that they will work over the term.

Content

Students will get an overview of applied machine learning with a focus on recent work in human-computer interaction. We will teach how to train a model using TensorFlow version 1.2 or later using Python 3.6. We will further focus on all nec-

essary steps to use a trained model within an Android application. This course covers machine learning models for two purposes, namely classification and regression tasks. Hence, the focus lies primarily on supervised learning which enables the models to be used for novel interaction techniques as shown in previous HCI work. For each topic, we provide multiple exercises that attendees will solve in the first part of the course.

In the second part of the course, students will apply the gained knowledge in practical group projects. In the projects, students will go through the whole development cycle:

- Exploration of different use cases
- Collecting data to train machine learning models
- Development of a neural network architecture
- Training the models using the collected data
- Deploying them on mobile devices
- Validating the deployed model

Learning goals

Students learn the basics of applied machine learning for interactive mobile systems. They are able to make informed decision when deciding for a machine learning model. They understand the process of using machine learning for interactive systems from collecting a data set, using it for training and validating the trained model.

Requirements

Basic knowledge of Python and good knowledge of programming is required to take part in the course. Students should have a basic understanding of machine learning and human-computer interaction.

Planned schedule

Week	Content
1. Week	Introduction to machine learning for human-computer interaction and the tools used in the course (e.g. Jupyter)
2. Week	Overview of machine learning and recent advances in the field, covering supervised & unsupervised learning, classification & regression, TensorFlow, a typical tool chain, and neural networks
3. Week	Participants train a neural network using provided Jupyter notebooks and explore the effect of different hyperparameters using a provided data set
4. Week	Discussion of performance metrics for classification and regression as well as cost functions. Overview of the process to bring TensorFlow models to Android devices
5. Week	Exploration of potential use cases

6. Week	Collecting data for the intended use case
7. Week	Designing a model that can be trained using the collected data
8. Week	Training the designed model
9. Week	Improving the model
10. Week	Deploying the model on mobile devices
11. Week	Testing and improving the deployed model
12. Week	Design of a study to validate the designed model
13. Week	Conducting a user study
14. Week	Evaluating the conducted study
15. Week	Presentation of the results

Expected Date & Time

Tuesdays, 14:00-15:30, SimTech 00.009 (Pfaffenwaldring 5a)
Available space: 6–14 students