



Prof. R. Wattenhofer

Advanced Topics in Deep Learning - Understanding and Evaluating Music Style Transfer

Style transfer is a recent and exciting topic in Deep Learning, and has so far been most successfully used in the image domain. However, there is also a lot of effort being put into performing style transfer in domains such as natural language and music. One of the biggest challenges is the evaluation. When was a style transfer successful, and how can we quantify success? One way is to perform listening tests with a large number of participants. Unfortunately, this method is time consuming and cannot be used to get feedback during the development phase.

In this thesis we want to gain a clearer understanding of the different possible types of style transfer for symbolic music (MIDI). We then want to investigate possible evaluation metrics to quantify the success of style transfers. For example, when defining style transfer as taking a Classic music piece and making it sound like Jazz, then a binary classifier could be applied before and after the style transfer. If the music piece is classified as Classic



before, and Jazz after the transfer, then we could say that the style transfer was a success. However, simple classifiers are usually not very robust and can easily be thrown off by, e.g., adding random noise to the input, which could reduce their validity as a metric. The thesis will involve lots of experiments to test different types of classifiers on different music datasets. What are the best features to predict the style of music? What is the influence of different pre-processing methods? Is it better to use CNNs or RNNs? Can we come up with different evaluation metrics altogether?

If this sounds interesting to you, do not hesitate to contact us.

Requirements: Knowledge in Deep Learning, or solid background in Machine Learning. Implementation experience is an advantage. You should be able to read and understand the first 12 chapters of the "Deep Learning Book" by Goodfellow et al. (available for free online from MIT press). If you are interested in the topic but new to deep learning we expect you to complete an introductory deep learning course before applying for the thesis, such as Andrew Ng's coursera course (use the free trial!)¹ or this Udacity course².

Interested? Please contact us for more details!

¹<https://www.coursera.org/specializations/deep-learning>

²<https://classroom.udacity.com/courses/ud730>

Contacts

- Gino Brunner: brunnegi@ethz.ch, ETZ G63
- Oliver Richter: richtero@ethz.ch, ETZ G63