ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Automatic Music Generation with Deep Learning

In this thesis we want to come up with new methods for automatic music generation based on deep learning. So far we have been focusing on symbolic music, i.e, MIDI files, which are similar to sheet music. However, we could also attempt generation of raw audio waves, which is generally more difficult and requires much longer training times.

Common problems of current symbolic music generation approaches are the limited number of simultaneous tracks/instruments and a lack of long-term structure. There are also limited possibilities for human input, e.g., we might want to generate music to suit a certain mood, or condition the generation on a "base" melody, or specify a subset of the instruments to be used, and so on. There are many interesting challenges to be solved on the road to better and more complete automatic music generation systems.

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 course[2].

Interested? Please contact us for more details!

Contacts

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

[1] `https://www.coursera.org/specializations/deep-learning`
[2] `https://classroom.udacity.com/courses/ud730`