



Prof. R. Wattenhofer

Generative Adversarial Networks for Music

Generative Adversarial Networks (GANs) are one of the most important recent advances in Deep Learning. Every week there are new papers using GANs in new creative ways and some of the results are truly amazing. There is even something called the GAN Zoo¹ containing more than 200 GAN specimen!

In the past we have had success in generating music and learning music theory concepts only from observing the data². We used some domain knowledge to construct an architecture that is able to capture long term musical structure. However, we used simple embeddings and LSTMs. Now we want to investigate using unsupervised generative models, such as GANs, to produce beautiful music. We have been working on this topic for some time and need your help to bring it to the next level! If this sounds interesting to you, do not hesitate to contact us so we can have a chat.



Requirements: Knowledge in Deep Learning (Equivalent to the Deep Learning lecture at ETH, or having read the Deep Learning Book³). Implementation experience with Keras and/or Tensorflow is an advantage.

Interested? Please contact us for more details!

Contacts

- Gino Brunner: brunnegi@ethz.ch, ETZ G63
- Yuyi Wang: yuwang@ethz.ch, ETZ 94

¹<https://github.com/hindupuravinash/the-gan-zoo>

²<https://arxiv.org/abs/1711.07682>

³<http://www.deeplearningbook.org/>