



Prof. R. Wattenhofer

Topics in Deep Generative Models, Disentangling and Style Transfer

Deep Generative Models have produced astounding content in the past few years. Especially the inception of Variational Autoencoders (VAE) and Generative Adversarial Networks (GAN) have accelerated the field and new papers are published every week. Generative models have allowed sophisticated style transfer between images (CycleGAN) and the creation of awesome music pieces (Google Magenta's PerformanceRNN).

Depending on your interests, we could improve on existing VAE or GAN architectures or apply the models to new tasks. Generative models have been used mostly for the visual domain, i.e., images and videos. However, there has recently been an increased interest into generating sequences, e.g., speech and music. A non-exhaustive list of possible topics is given below. However, you are welcome to come up with your own ideas. If you are interested in deep generative models, do not hesitate to contact us! We will then provide you with a more detailed list of current topics.



- Musical style transfer with raw audio or MIDI
- Speaker style transfer for text to speech
- Style/sentiment transfer for text
- Learning disentangled latent representations with VAEs or GANs (for images, videos, music, language, ...)
- Applying GANs for sequence tasks (language, music, ...)
- Improve (Conditional) CycleGAN for more fine-grained control
- Speech reconstruction/superresolution with GANs
- Improving VAEs or GANs, performing experiments, comparisons, etc.
- ...

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.

Contacts

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

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

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