

Generating Art with GANs

Generative Deep Learning has received considerable attention in the last couple of years. This has sparked repeated interest in an old, philosophical question:

Can artificial intelligence produce art?

Or, more general:

Can machines be creative?

Deep learning research has given a practical answer by producing seemingly creative artefacts: composing music,¹ writing novels,² or screenplays.³ In addition, there have been notable advances in using AI for the visual arts. Two years ago, Christie's auctioned off a portrait-like painting generated by a neural network trained on 15,000 portraits painted between the 14th and the 20th century (see article on the right). In contrast to music and text, where sequence models such as LSTMs and Transformers are primarily used to generate the output, visual art is mostly generated using Generative Adversarial Networks (GANs).



<https://www.christies.com/features/A-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx>

Project Outline

In this project we want to go beyond simply training a GAN to create visual art. Our goal is to influence the “creative” process by providing external knowledge and emotional cues, e.g., “Create a *SAD* painting of a *DOG*”. To this end, we split the project into three phases:

1. Learn mappings of emotions to paintings,
2. Learn mappings of objects to paintings, and
3. Generate paintings given emotion and object.

In the first stage, we focus on emotions and try to generate paintings given one emotion or a combination of different emotions, e.g., happy + surprised. In the second stage we try to generate paintings given one topic/object. Optionally, we could also try to generate combinations of objects, e.g., a bridge + a ship. In the third stage, we want to be able to generate paintings with a certain emotion and a certain topic.


¹ <https://www.youtube.com/watch?v=Emidxpkyk6o>

² <https://medium.com/the-research-nest/interesting-novels-written-by-artificial-intelligence-d407e330fe07>

³ https://www.youtube.com/watch?v=5qPgG98_CQ8

Project Approach

We will learn the mappings using annotated/labeled training data. A new dataset called Wikiart Emotions (Saif Mohammad, 2018) contains paintings labeled by humans with associated emotions. We can also use general image datasets that do not focus on art. The popular image database ImageNet (Jia Deng, 2009) contains 15 million images with 20k different objects. We will explore the possibility of using this data directly or if needed apply neural style transfer (Yongcheng Jing, 2019) to create annotated paintings.



We will adapt the standard GAN architecture to incorporate information about emotions and objects. Conditional GANs have been successful in the past for similar tasks (Bo Dai, 2017). We will develop a similar approach to solve our task.

We are planning to write a research paper at the end of the project for submission to a conference. If you are interested in doing research in deep learning at the intersection of

- Computer vision,
- Text mining, and
- Digital cultural heritage,

then this is the right project for you. It would be helpful if you have some prior knowledge about deep learning architectures and deep learning development frameworks such as PyTorch.

Contact

This project will be jointly supervised by the Information Systems chair and the AI and Intelligent Systems chair. If you have any questions, please do not hesitate to contact Dr. Krestel (ralf.krestel@hpi.de) or Prof. Dr. de Melo (Gerard.DeMelo@hpi.de).

References

- Jia Deng, e. a. (2009). Imagenet: A large-scale hierarchical image database. *IEEE conference on computer vision and pattern recognition (CVPR)*.
- Saif Mohammad, S. K. (2018). Wikiart emotions: An annotated dataset of emotions evoked by art. *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC)*.
- Yongcheng Jing, e. a. (2019). Neural style transfer: A review. *IEEE transactions on visualization and computer graphics*.