

ADGN20: First workshop on Applied Deep Generative Networks

Jer Hayes*, Cathal Gurrin●, Azzurra Pini°, and Mark Keane

*jeremiah.hayes@accenture.com, ●cgurrin@computing.dcu.ie,
°azzurra.pini@jordnet.com and mark.keane@ucd.ie
<https://sites.google.com/view/adgn-20/>

Abstract. Generative models are widely used in many subfields of AI and Machine Learning. More recently generative models using deep learning have been employed in a creative manner to generate new media (images, text and music), but they have also been applied to areas such as drug discovery and data anonymisation. This workshop focused on the applications and research problems related to the practical use of deep generative models in the real world and where examples ranged from drug discovery to bespoke furniture generation.

Keywords: Generative Networks · GANs · Applied Research

1 Background

A generative model is a powerful way of learning data distributions using unsupervised learning; essentially the model learns unseen patterns or hidden structures in data in order to synthesis new example data. When we speak of deep generative networks we are describing models that are built as deep neural network (or deep learning) models and that typically are grouped into types of autoencoders (especially variational autoencoders [1]); generative adversarial networks (GANs) [2]; and combinations of the two. The main applications of generative networks have been based on using image data. Given an image dataset new example images are created, e.g., [3]. Depending on the dataset these can have specific domains, e.g., generating realistic human faces [4]; anime characters [5]. Another use case is ‘translation’ of one input into another. These include transforming satellite photographs to Google maps images, transforming sketch images to realist color photographs. A significant number of image-to-image translations are called ‘style transfers’ and they include examples such transforming photos from day to night; transforming black and white photographs to color; transforming a painting to a photograph; and transforming a photograph to a painting (with a certain artistic style). However, there have also been examples where text is used in association with images, e.g., text-to-image translation where an image is created based on the text input. Some of these techniques have also had negative applications, e.g., the use of ‘deepfakes’ [6] whereby realistic fake content is generated sometimes for purpose of creating fake news content.

The intention of the workshop was to focus on research problems related to the practical use of deep generative models in the real world. It was also intended to be an in-person workshop but travel restrictions due to the Covid-19 pandemic resulted the workshop and indeed the whole conference being held virtually. Ultimately seven papers were accepted for presentation at the workshop. These papers explore topics mainly related the following:

- Business / Service Applications
- Exploration of latent space
- Human-machine collaboration

The papers were divided into sessions that covered business and health Applications; human-machine collaboration; and the latent Space. We now briefly introduce these seven contributions.

2 Work presented at the workshop

Relating to business / health Applications, Pandey et al., apply GANs to the general problem of dealing with data that is imbalanced. This occurs in many sectors but it is especially prevalent in the financial sector where legitimate transactions make up the vast majority of total transactions, thereby making it harder for to detect fraudulent transactions. They presented a study on GANs for synthetic fraud data generation and demonstrate improved classifier performance for detecting fraud. In related work, Zhang and Li’s paper (they could not attend the virtual workshop) introduced research on generating latent weights for few shot image classification. This work focused on another common problem building models from limited labeled data. Additionally Gaàl et al. presented research on adding an adversarial critic model to a CNN in the application area of Chest X-ray Lung Segmentation where both performance and interpretability of the model results are important.

Within the domain of human + machine collaboration, Pini presented work that takes a design approach to style transfer. This work is an application of season transfer with GANs to design a visual booking service for the travel industry. González and Muiños-Landin Santiago then presented work on generative design for social manufacturing with a use case featuring furniture design. Given the remote nature of the workshop we unfortunately could not hand the physical outputs of their use case.

Research into exploration of latent space included Singh et al. who presented work on manifold traversal of latent spaces for novel molecule discovery. Fernandes et al. also presented work on latent Space exploration for classifier improvement. They proposed a framework that combined GANS and evolutionary computation to perform data augmentation on small datasets in order to improve the performance of image classifiers trained via supervised learning.

3 Conclusion

Despite the vagaries of remote meetings the workshop was held as originally intended and provided a form for both presentation and discussion. Six of the seven accepted papers were presented during the workshop. In terms of the discussions that took place the focus was on the move towards real-world application and the implications for ethics, collaboration, and interpretability of bringing GANs from the lab into the real-world. It was highlighted that considering technology application was opening up new research opportunities and that this move to application-driven research was mirroring the situation in other fields of AI and machine learning. Collaboration was a talking point as currently many of these models and techniques are used by experts but Pini pointed out that products like RunwayML (<https://runwayml.com/>) allow non-ML experts to use generative networks. The question of interpretability also arose as currently for some models it is not clear as to why a particular example was synthesised etc. Overall the discussion was robust and involved the organisers and many of the participants. The organisers noted the potential for follow-on workshops that will facilitate a continuation of the discussions.

References

1. D. P. Kingma and M. Welling, “An introduction to variational autoencoders,” *Foundations and Trends in Machine Learning*, vol. 12, no. 4, p. 307–392, 2019.
2. I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial networks,” 2014.
3. A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” 2015.
4. T. Karras, T. Aila, S. Laine, and J. Lehtinen, “Progressive growing of gans for improved quality, stability, and variation,” 2017.
5. Y. Jin, J. Zhang, M. Li, Y. Tian, H. Zhu, and Z. Fang, “Towards the automatic anime characters creation with generative adversarial networks,” 2017.
6. R. Tolosana, R. Vera-Rodriguez, J. Fierrez, A. Morales, and J. Ortega-Garcia, “Deepfakes and beyond: A survey of face manipulation and fake detection,” 2020.