Generative Adversial Networks (or GANs) are a new method used in AI to generate synthetic data. The first GAN paper appeared in 2014, and since then, the number of applications has grown rapidly.

Perhaps the most famous example is that of facial image generation. The idea is to train a model to recognise common features in faces. This is done the same way that neural networks are often trained: by showing the network many thousands of images of faces, and the network's job is to identify the common features, as well as where they appear in relation to each other (e.g. the mouth should be below the nose, the two eyes should be next to each other etc).

Below are some examples of faces that have been generated by a trained GAN (obtained from https://generated.photos/faces). These faces are completely synthetic, i.e. they are not real people at all, merely faces that include the features that are commonly found in a human face.

If I hadn't told you, you probably would not have guessed that none of these people exist. This is an incredible achievement considering that computers don't actually understand the task they are performing. But what does it have to do with my interests?

As is often the case with AI methods, there are many other possible applications of GANs besides the single example given above. As well as images, GANs can be used to learn the features of time series data, and to then make predictions about how that data will appear in the future (think stock price predictions). One nice example of this comes from music, where the sound produced by an instrument can be recorded as a time varying signal.

Below is an example of a real audio waveform, and different signals generated by GAN models that first learned the features of a training set of audio samples, and then used that 'knowledge' to generate new, synthetic sound waves.

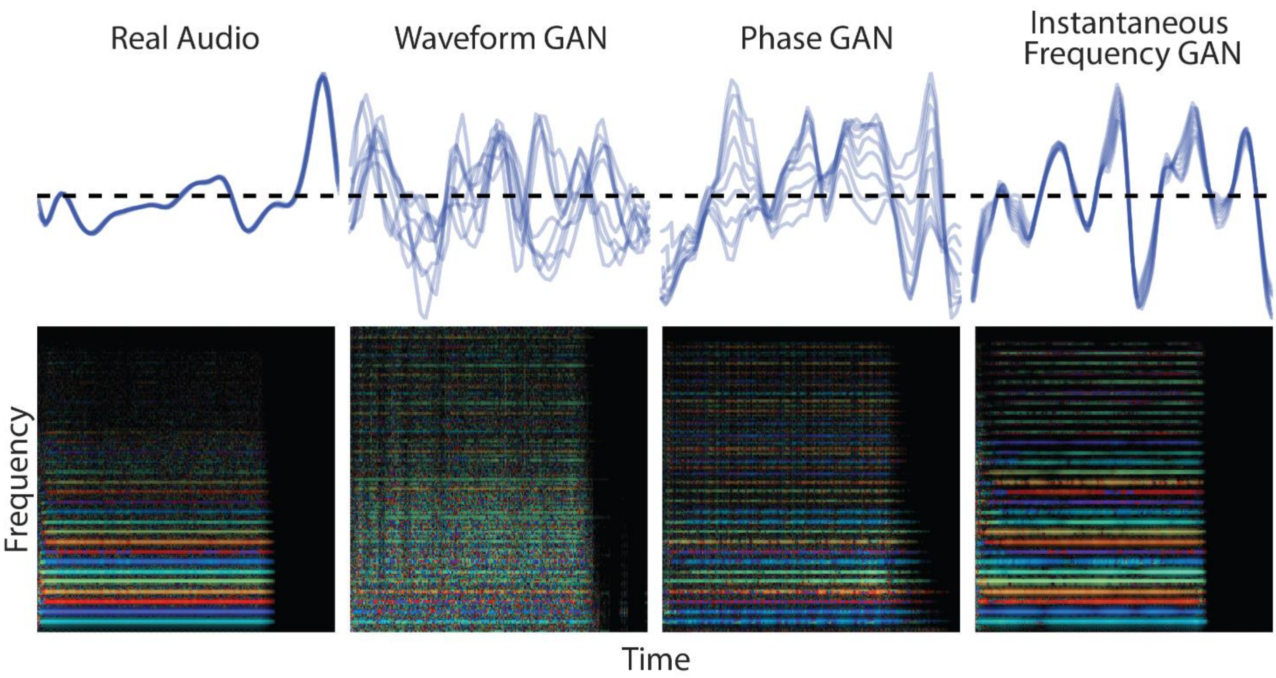

Image from https://magenta.tensorflow.org/gansynth

I played around with a similar method called museGAN, and it does seem to be a very promising way to learn the features of this kind of signal, and to use that information to generate your own samples (listen to my ‘music’ here: gan_music.mp3).

As interested as I am in music, it is not what I get paid for (probably fair based on my musical GAN efforts). I had been considering possible applications of GANs for some time before I realised that one important use case for this kind of approach is to generate large volumes of data that can be used to train deep learning models. I wondered whether this could be used to help me increase my ultrasound image dataset, which could in turn be useful for training better models for the automation of data analysis. As deep learning methods often require large labelled datasets, I wanted a way to automate the process of both synthesising **and** analysing ultrasound images, meaning that the user would not have to actually manually label anything. In other words, I wanted to develop a fully unsupervised approach.

For this I turned to a paper that used GANs to generate images, as well as to automatically segment the images (I focus on the first part here). I used a set of real ultrasound images and a set of synthetic masks (see the original paper for methods details), and then trained a GAN to produce ultrasound images that could conceivably be mistaken for real images. Our manuscript describing this work is now published here, but below is a set of images, some of which are real and some of which I have generated using my GAN approach. See if you can guess which ones are which (scroll down to the 4th page to find the answers).

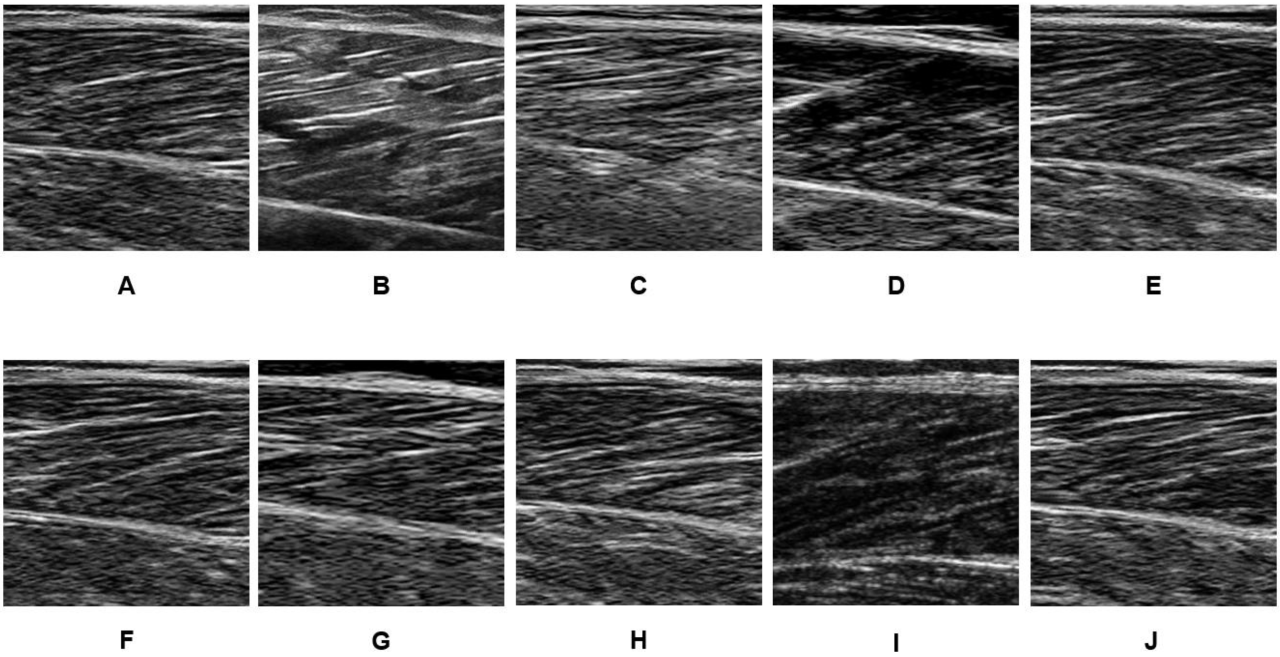
A
B
C
D
E
F
G
H
I
J

Real: B, C, D, G, I
Synthetic: A, E, F, H, J