DAVID RANNALEET & VICTOR GUNNARSSON

Generating brain signals by taking inspiration from diffusion models

Just as how *Stable Diffusion* can generate thousands of non-existing images of dogs, can we use the same method to generate non-existing brain signals? We found that not only can we use the same method, but it can also be used to improve decoding of real brain signals.

A 'Cocktail Party' Environment

Most of us have experienced the struggle of following a conversation at a noisy party, family gathering or maybe a crowded train station. Despite this, our ears and brains remarkably manage to focus on one particular conversation, effectively tuning out the unwanted noise. The problem of figuring out how our brains do this is known to researchers as *the cocktail party problem*. What would be the benefit of knowing how the brain does this? Well, we know that this is one area where the hearing impaired struggle quite a bit more than the rest of us. Many hearing aid users often remove their devices in noisy environments, as these aids typically do not match the brain's ability to single out and amplify the intended speaker. Imagine if these hearing aids could tap into and 'decode' the wearer's brain activity to discern their focus of attention. Is such a thing possible?

Brain Activity to Electroencephalography

Electroencephalography (EEG) provides that possibility, offering a non-invasive glimpse into brain activity by recording from electrodes placed on the scalp. In our work, we used EEG to discern whether individuals, particularly those with hearing impairments, were focusing on a speaker to their left or right in a simulated noisy environment. Since it is quite cumbersome and expensive to collect this EEG, we wondered: Could we create new, synthetic EEG samples from existing ones?

The Magic of Diffusion Models

Enter diffusion models, which have seen widespread success in products such as *DALL-E* or *Stable Diffusion*, which generates thousands of non-existent images of fantastical dogs, landscapes, or even astronauts on horses! Their success in crafting intricate visuals suggested diffusion models might handle the complexities of EEG data too. In layman's terms, the model works by starting from some pure noise, such as a TV static screen that many of us have might have seen. From this pure noise, the model learns to remove the noise gradually until you are left with an image, say, of a dog.

Exciting Results

Our experiments with many diffusion models ended up being successful, and we produced synthetic EEG that closely matched the real thing. With the EEG looking good, we then added this new generated data to the algorithm that was responsible for classifying whether a person was trying to listen to the speaker on the 'left' or 'right'. The results? We saw an average increase in classification accuracy of **0.58%**, improving to almost **71%** from the basic model result of **70.4%**. While on the surface this might look small, it's remarkable when considering that we achieved this by treating complex EEG signals as images! It's very interesting that the diffusion models can produce high-quality EEG data without the need for the complex manipulations one might think are necessary for such signals. We leave it to future research to now optimize the generation process to be more tailored for EEG, since we believe there is a lot of room for improvement. ■