Project Description:

We aim to use deep learning / reinforcement learning techniques to understand the non linear distortion to incoming signals caused by thermodynamics effect in power amplifiers. Power amplifiers are a critical component in communication systems. Previous research have delved into this area by creating mathematical and statistical models to understand this non linear distortion, and the standard the standard technique is to use a pre-distortion module before the power amplifier which can reverser (invert) the non linear effect of the power amplifier. Our aim is to improve this technique by using deep learning/ reinforcement learning, and make it so that our proposed technique can be used to linearize a wide range of signals, and provide theoretical guarantees to our techniques, and finally potentially turn this into a commercially available product (by making the process faster and more efficient).

Project Phase I Description:

Phase I is a proof of concept process, namely, our team investigate whether we can build a simple, small, invertible neural network to capture the non-linear distortion caused by a simulated power amplifier (by a memoryless Saleh model or a Saleh model) to a simulated OFDM or Gaussian signals with phase difference. Since this phase has started a few months before I joined this project, I think it took us at least 6 months to complete this and turn into an academic paper. In the end, our team has successfully build an invertible recurrent network using customized activation function and leverages the idea of normalizing flow. The entire network has less than one thousand parameters, which is atypical in large-scale deep learning. This point is very critical to our goal since our ultimate aim to build a fast and efficient way to linearize signals. To my understanding, Phase I will be completed in mid September. Since our proposed network has only worked in simulated signals, we imagine it would be quite a lot of work when we try to expose the network to more complex signals, and the architecture of the network will need to potentially change a lot, this will be Phase II.

My role in Phase I.

Since signal processing and commutation theory is a very new topic to me, I spent the majority of my time learning a lot of the background knowledge, running experiments and maintaining our code base. For signal processing, specifically I have brushed up on my Fourier analysis; I have learned the basics of signals and systems; the Nyquist sampling theorem; frequency response and analog filter design. For communication theory I have learned about path loss, shadowing, multi path channel models, channel capacity, channel coding theory, modulation theory (digital modulation and adaptive modulation, multi carrier modulation that includes the Orthogonal Frequency-division multiplexing signals and how to generate them), multiple antennas space-time communication theory, etc. These materials are about 2 years of wireless communication graduate-school study, and obviously I have only understood the basics of them within my short time studying them, it was really fascinating though. For deep learning I have brushed up on my pytorch, recurrent neural network, non-activation theory (its impact on the optimization terrain of the neural network using a stochastic gradient descent method), normalizing flow, etc. Some of the more important things about deep learning I have learned while running experiments for my team includes how to monitor the overfitting of the model, how to reduce overfitting by merging not active layers together while reducing parameters. I have also learned about the more practical side of signal processing, such as using Matlab to generate OFDM signals, applying filters, and I have implemented several important metrics to measure the degree of linearization we have achieved for our signals, such as cutoff probability. I have also spent a lot of time refactoring our code base, while doing so I think I gained a deeper understanding of our model and its performance. I have also spent a considerable amount of time studying reinforcement learning using Richard Sutton’s book, which is also a new area to me. Reinforcement learning is also another technique that we hope to apply to sampling. Specifically, I have learned about the basic concept of reinforcement learning, multiple on-policy and off-policy algorithm such as monte carlo control, TD-lambda learning, SARSA, Q-learning, double Q-learning. I have spent time researching the intersection of deep learning and reinforcement learning which is the hottest field right now ever since the advent of large scale neural networks. I also spent a bit time learning about the convergence guarantees for each of these algorithm.

Overall, I think I have learned a lot during the summer, and I am confident that I will be able to take a more active role in contributing to the second phase of our project. Our team is still meeting to discuss the specifics of the second phase of our project, which I will probably be able to share them