Course: EE P 518 **Title: Signal Processing: Deep and Classical Methods** **Credits: 4 Units**

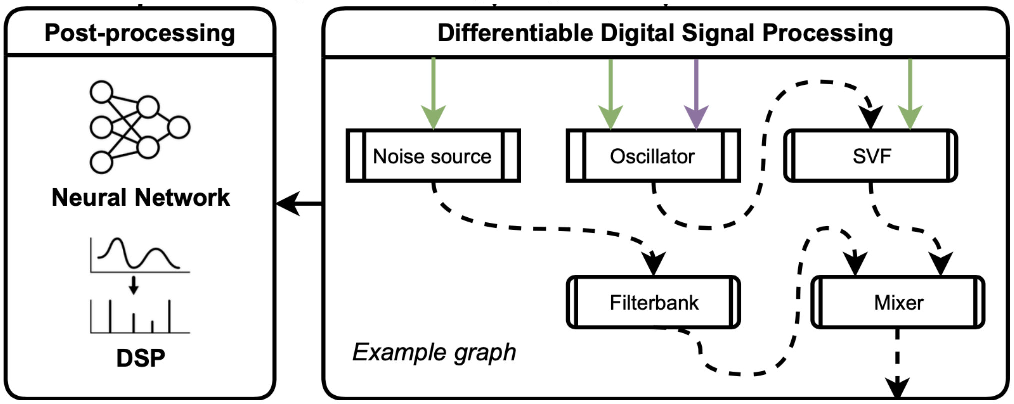

Lecture and Discussion Time: Tuesdays 6:00–8:50 pm in ECE Building Room 037

Why take this class? This class is intended for engineers and scientists within and outside Electrical Engineering. What is it useful for? Some examples, from basic to advanced: How many engineers and scientists know that linear interpolation is a poor way to increase sample rate? What are the best signal representations for machine learning systems which generalize well, that is do well outside their training data? How do linear time invariant systems generalize? And how do modern deep nets go beyond the capabilities of linear systems?

Over the last 45 years digital signal processing has gone from classical to deep methods. The term "differentiable digital signal processing" describes a new family of techniques in which loss function gradients are backpropagated through digital signal processors, facilitating their integration into neural networks. We will discuss recent results and include such investigations in our final lab project.

Classically, why is a Fourier transform and frequency important? How does the concept of frequency have principled depth which goes way beyond simply decomposing arbitrary signals into oscillating components? What are the limitations of Fourier transforms, and how do more general deep networks get around these limitations?

A fast Fourier transform is a well-known fast implementation of a discrete Fourier transform. How does the fast Fourier transform potentially massively speed up the discrete Fourier transform? But how is a discrete Fourier transform (and hence a fast Fourier transform) often a really poor approximation of a true discrete-time Fourier transform? Why is a fast Fourier transform usually unsuitable for signal filtering?

Also: How can you best characterize signals from time-varying systems, like real-world examples of speech, patches of images, and video? Why do Fourier transforms not fit this problem and how can they be modified? How do adaptive, Wiener, and Kalman filters fit this problem? Why are somewhat more recent and popular concepts like convolutional[1] and recurrent[2] neural and deep networks potentially more than simply hype and perhaps useful for your future engineering work?

EE P 518 Prerequisite: *Recent undergraduate signal processing or other quantitative, particularly Fourier transform, theory background is required. You need to be comfortable with complex numbers, know at least elementary matrix theory, and know what a Fourier transform is.* If you have an undergraduate background in DSP, this class will still be challenging. For example, it will cover graduate level DSP concepts such as signal processing for signals from time-varying systems, multirate signal processing, and non-Euclidean decomposition spaces, which are needed to understand much of the more advanced signal processing, control, and related literature.

This EE 518 course quickly reviews linear time-invariant systems, discrete-time signals, sampling, Fourier transforms and bilateral *z*-transforms. Note that some on-line courses also provide a fine background, but most don't provide the kind of in-depth problem solving or recent results that will be central to EE P 518.

---

[1] For first use of convolution within artificial neural networks, see right side of figure 1 of Homma, Atlas, and Marks, "An Artificial Neural Network for Spatio-Temporal Bipolar Patterns," *Proc. Neural Information Processing Systems (NIPS)*, 1987.

[2] For an early paper on recurrent networks, which also overlaps statistical signal processing, see: Connor, Martin, and Atlas.