EE 367 Project Proposal

Leo Dong

February 18 2022

Motivation

Real world signals are conventionally captured and recorded by sensors as discrete-valued functions. For example, a digital image is a discrete-valued function that is piecewise constant at each pixel coordinate. But this is not a faithful representation of the physical world, which is mostly continuous. An alternative representation is using parameterized implicit functions. To use images as an example, we can define $f_\theta(x) : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ with parameter θ, such that $f_\theta(x)$ maps any (x_1, x_2) coordinate on a plane to an RGB value. f is typically represented with a neural network with parameter θ.

SIREN is a particular type of neural networks that uses sinusoidal functions as activation functions. It has been demonstrated that SIREN is particularly suitable for learning implicit functions for signal representation, because it preserves spatial details and derivatives information from the input signal. In this project, I would like to experiment using SIREN to solve the deconvolution inverse problem in a self-supervised setting. In other words, given that SIREN is a good signal representation method, I want to see whether it is robust to signal noise when learning its function parameters.

Planned Experiments

I will first implement SIREN using the same architecture that the original paper used for its image signal representation experiments. In order to make SIREN perform deconvolution tasks, I will regularize the SIREN reconstruction using the TV norm in addition to its data fidelity term. This will give a good idea of how easy it is to separate true signals from noise in a SIREN network.

Next, I want to see if SIREN's inherent architecture using periodic functions as activations is a good prior for natural images. This will be a similar experiment to the Deep Image Prior paper that analyzes convolutional neural networks. Specifically, I will train SIREN to overfit to a single image with noise, and see if it is able to first reproduce a clean image before overfitting to the noise.

If I still have time left, the last experiment I will try is to see if I can incorporate SIREN into the ADMM optimization framework. The most straightforward way is to use the denoiser I learned in the first experiment, althought it might not be very interesting since that denoiser was learned in a self-supervised fashion using TV norm in the first place. If the second experiment is promising, a more interesting experiment is to see if I can recover the implicit prior in SIREN networks in some smart formulation that can fit into the ADMM optimization pipeline.

Timeline

By end of Week 8: Set up SIREN network that is able to reproduce the original paper's experiments.

By end of Week 9: Prepare small image dataset and manually add noise to it. The dataset does not need to be too large because I am not relying on learning image priors from data. Run initial denoising experiments with TV regularization.

By end of Week 10: Run Deep Image Prior type of experiment on SIREN. If have time, see if I can incorporate SIREN into the ADMM optimization pipeline.

Related Works

Before SIREN, there are many works in designing implicit functions for representing 3D shapes, such as Occupancy Networks, Deep Signed Distance Function, and Local Deep Implicit Functions (LDIF). Notably, Local Deep Implicit Functions exhibited particularly high fidelity reconstruction results because it learned each part of a 3D shape separately. Inspired by this locality-based approach, Local Implicit Image Functions and Local Functional Representations are two recent methods for representing images specifically.

Even though locality-based approaches demonstrate high reconstruction fidelity, they usually come at a cost of over-parameterization and increased network or training complexity. Therefore, I want to focus on SIREN itself and the property that is brought by this simple yet powerful idea of using periodic activation functions, without diving too deep into the heavier engineering of the highest fidelity locality-based methods.