Image Denoising in Real Analysis

Possible Courses: A one semester course in Real Analysis.

Application Suppose you have an image that has been corrupted by noise. What features of the image tell you that it is noisy? How might you produce a cleaner image? Can you always (ever?) recover the original image?

Motivated Concepts This module motivates the study of metric spaces, Cauchy sequences, and derivatives. Students see that understanding derivatives can help in solving real problems that they care about.

Prerequisite Material This module is designed to be introduced with no prerequisite knowledge. For later labs in this module we do assume that students have some familiarity with topics from multi-variable calculus (though we don't assume fluency - and actually we expect that the module will help to reinforce/reacquaint our students with these topics). The labs will be introduced throughout the semester, and we provide recommendations for timing (and what should and should not be covered before each lab) in the lab descriptions below.

Description Images can be represented as vectors of pixel intensities. Various metrics give distinct ways to quantify similarity between images. Students will be able to explore how different metrics might be useful for different situations.

If d is the original image, what is a candidate for a denoised image? To answer this questions we note that there are two competing goals in denoising: one is to remove noise, and the other is to preserve information. (The blank image has no noise, but also contains no useful information!)

One possible denoising solution is the image u that minimizes the quantity $L(u) := \sum |D(u_i)| + \sum |u_i - d_i|$. (Here the quantity $D(u_i)$ is a discrete gradient approximation at the i-th pixel.) Many variations of $L(u)$ might be used here (for example, weighting one of the terms in the sum more heavily), and a major part of this project is for students to see that the different variations give rise to denoising solutions with differing properties.

Once a variant $L(u)$ (the denoising functional) has been determined, we then turn to the problem of actually finding the minimizer. Given a starting image u_0, we can calculate the value $L(u_0)$, but how do we then produce an image u_1 with a lower value $L(u_1)$? In this case we draw on the students' intuition from multivariable calculus. (And indeed the function L is a map from a finite dimensional space to $\mathbb{R}$, so they are in essentially the same situation here.) Note that *directions* in the space of images are themselves images, so we want to choose a direction image and add a multiple of it to u_0. The direction image to choose is exactly the one that corresponds to the gradient of the function $L(u)$. The step size corresponds to the multiple of the direction image that is added to u_0. An optimal step size can be found using directional derivatives, and once a direction has been fixed, these directional derivatives are just the ordinary one-dimensional derivatives from the course!

The project opens up a whole host of new questions for students to consider. The process of finding the final denoised image is an iterative one, so it is important to know whether the algorithm converges, and how fast it converges. Also, even if it converges, must it converge to the original clean image? (Spoiler: it need not, and usually won't.) Does the choice of metric affect the denoising solution? There are also questions regarding the space of images. For example, is the space compact? (Does this depend on the metric being used?) There are very rich connections to functional analysis that can be explored.

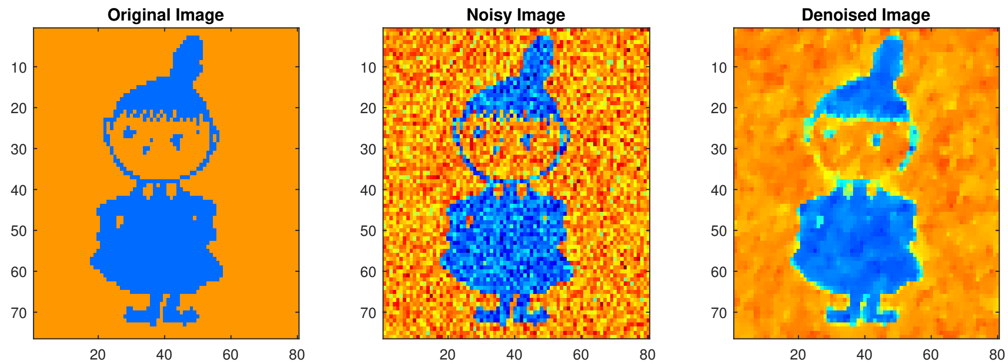

Figure 1: Examples of images students might produce in the module.

Time Frame The module is designed to be interleaved in a one-semester real analysis course.

Lab 0 *What is noise?* In this exercise, students are given a noisy signal and asked to draw, by hand, a denoised version, taking into consideration what appears to be noise. In the second part of this lab, students use the idea of a metric to discuss how to measure noise in a signal. (Can be given as homework.)

Lab 1 In this lab, students will be given noisy images and the corresponding clean image. They are asked to list the properties of noise in an image. Images are defined as vectors in a finite-dimensional vector space that can also be thought of as a metric space. Students are asked to explore different metrics on the set of images, with the goal of measuring the distance between a noisy and a "denoised" image. (1-2 classes)

Lab 2 *Creating the Denoising Function.* In this lab, for a given noisy image d, students create a real-valued function on the (finite-dimensional!) space of images that measures how good an image is as a denoising of the data d. (Of course, the function depends on d.) (1-2 classes)

Lab 3 *Sequences of Images.* In this lab, students are given part of a sequence of images created in a denoising process. Students are asked to discuss convergence of sequences based on successive terms getting closer together. This lab leads to the need for the Cauchy Criterion. (1 class)

Lab 4 *Directional Derivatives.* Given a direction to move in the vector space of images, students consider how the denoising changes as they move in that direction. For a fixed direction this is a 1-dimensional problem, and with this in mind students consider the infinitesimal change in the denoising function as a derivative. This lab can be used to motivate the study of derivatives. (1-2 classes)

Lab 5 *Gradient Search.* In this lab, students explore the method of gradient descent in order to minimize the function that they created in Lab 2. Subtle issues (too large of a step size or shrinking step size too quickly) that can cause the algorithm to fail to reach a minimizer are also explored. (1-2 classes)

Lab 6 *Denoising Challenge.* In this lab, students are asked to summarize the image denoising process they have developed. Students are then asked to provide an image in which noise will be added. The students will have a final task of using Matlab code to denoise their (corrupted) image by adjusting the parameters. The lab requires a students to explore various choices of parameters and how they change the "denoised" image, and to articulate why and how they arrived at their solution. An optional (but highly encouraged) assignment included in the lab materials asks students to incorporate their solution into a scientific paper where they explain the problem, discuss the mathematics they have learned, and explain their solutions. (Most can be given as an out of class assignment, so anywhere from 1-4 class periods can be used for this part of the module.)