Image Denoising in Real Analysis

Possible Courses: A two-semester Real Analysis sequence.

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 (1) metric spaces and in particular function spaces, (2) mathematical modeling and gradients, (3) iterative processes, sequences and Cauchy sequences, and (4) gradient search methods and ideas from the calculus of variations.

Prerequisite Material Students should have already covered the following material in Real Analysis before beginning the first day of the module:

1. Metric Spaces (supplementary materials are provided for this topic)

The class **should not** have covered the following before the module:

1. Sequences in Metric Spaces
2. Sequence Convergence
3. Calculus of Variations

Individual prerequisites for each lab are described in the lab documentation.

Description Images can be represented as functions on a discrete set of pixel locations. Various metrics (e.g., L^1, L^2, L^∞) give distinct ways to quantify similarity between images. Students will be able to explore how different metrics might be useful for different situations.

If f is the original image (function), 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 (function) u that minimizes some quantity $F[u]$, where a lower value indicates better quality or less noise. One example is

$$F[u] := \int |\nabla u(x)| + \int |u(x) - f(x)| \, dx,$$

though any number of variants might be used (such as weighting one of the terms more heavily). A major part of this project is for students to see that the different variants give rise to denoising solutions with differing properties.

Once we choose a variant, (calling it $F[u]$) to serve as our the denoising functional, we then turn to the problem of actually finding the minimizer. Given a starting image u_0, we can calculate the value $F[u_0]$, but how do we then produce an image u_1 with a lower value $F[u_1]$? In this case we draw on the students' intuition from multivariable calculus. 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. Students investigate topics from the calculus of variations to discover that a fruitful direction image to choose is one that corresponds to something like a gradient of the functional $F[u]$. Students explore how important it can be to choose an ideal step size, but to also allow the step size to change. Students explore directional derivatives by considering the one-dimensional parameterization in the direction of descent. Applying their understanding of 1-D minimization, students iteratively find a minimizer and then a new direction of descent until a stationary point is found.

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.

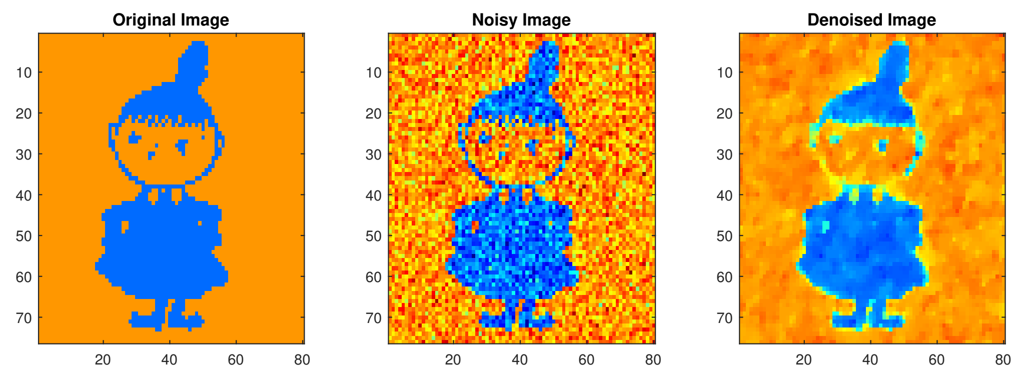

Figure 1: Examples of images students might produce in the module.

Time Frame The first part of this module (the Pre-Lab and first 2 labs) are intended for use in a first semester real analysis class. The first semester material will take 2-3 class periods: 1-2 near the beginning and the last one in the middle (near the discussion of sequence convergence). The second part (the remaining labs) are intended for use in the second semester. The second semester material takes an estimated 2 weeks of class time, although it is likely that instructors will also want to supplement their courses with some "extra" material that is not necessarily standard for the second semester course. Some such topics include function spaces, linear functions, and gradients.

- [PreLab] Can be assigned as homework. In this lab, the 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.

- [Lab 1] Students are given a set of noisy images and the corresponding clean images. They are asked to list the properties of noise in an image. Then students are asked to consider the set of images as a metric space, with the plan to measure the distance between a noisy and a "denoised" image. They develop several different metrics and explore the differences between those metrics.
  An important component of this lab is to define an image, and to emphasize the distinction between discrete and continuous functions.
  (1-2 classes)
- [Lab 2] 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 3] This lab is the first lab of the second semester course. In this lab, students are reminded of the properties of noise in an image. Students see how to represent images by functions defined on a subset of $\mathbb{R}^2$. They explore different ways to measure variation in an image (and within in the function representing the image) and measure data fidelity: how close an image is to the data given (and the distance between the functions representing the data and a denoised image). This is the lab in which students develop the denoising functional. (Roughly 2 classes)
- [Lab 4] This lab starts the process of actually finding a candidate denoised image. In the lab students use Matlab code as they think about what a "direction" in image space is. They then look at what happens to the denoising functional from the previous lab when they change an image by adding a multiple of a direction image. (One class period.)
- [Lab 5] In this lab, students explore the method of gradient descent in order to understand minimization of a functional using the Euler-Lagrange equations. 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 class period)
- [Lab 6] In this lab, students are asked to discuss how the calculus of variations approach makes sense in the task of denoising an image. Students are then asked to provide an image in which noise will be added. The students will have a final task of denoising their noisy 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. This lab works best as a final project that 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. (This lab can be given as an out of class assignment, so no more than 1 class period is likely to be used for this part of the module.)
- [COV Discussion] Just as 1-dimensional derivatives can be used to find (local) minima of functions, variational derivatives can be used to find (local) minima of functionals. This interactive class discussion/assignment will guide students through the development of variations (variational derivatives) of a functional through 1-dimensional derivatives and the multivariable chain rule. Students will then use the variation to find the Euler-Lagrange equation of a general functional and then the denoising functional. They will the Euler-Lagrange equation to find a direction of descent, in the

function space, for the denoising functional. (1-2 class periods, possibly with homework assigned.)