Image Processing in Analysis 1
Lab #2

Group discussion:

Recall that the set of $m \times n$ images can be thought of as a metric space and that there are many choices of metrics (or distance functions) on this set.

Consider the following images. The first image contains noise, while the second image is the true (starting) image.

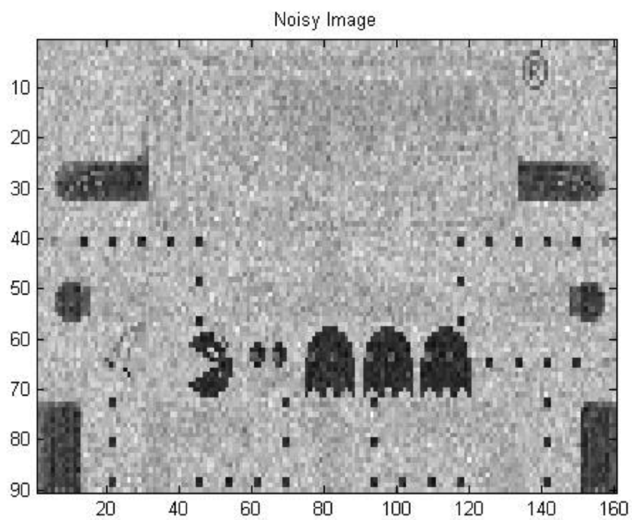

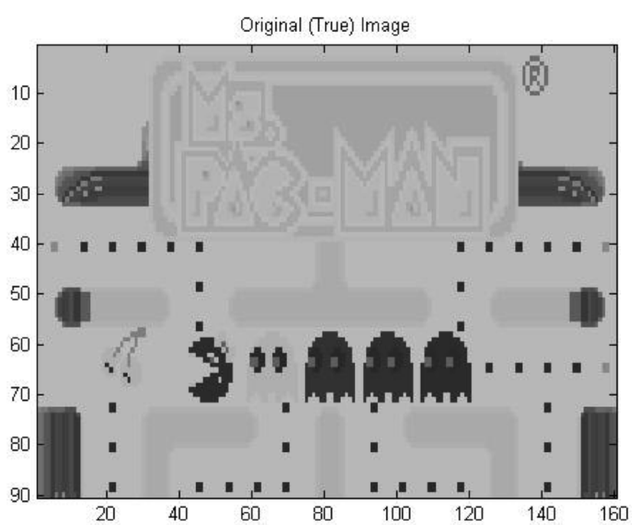

An important task is to take a noisy image and produce a new "denoised" image that looks as much like the original as possible.

The figure on the next page shows the intermediate images produced by a denoising algorithm.

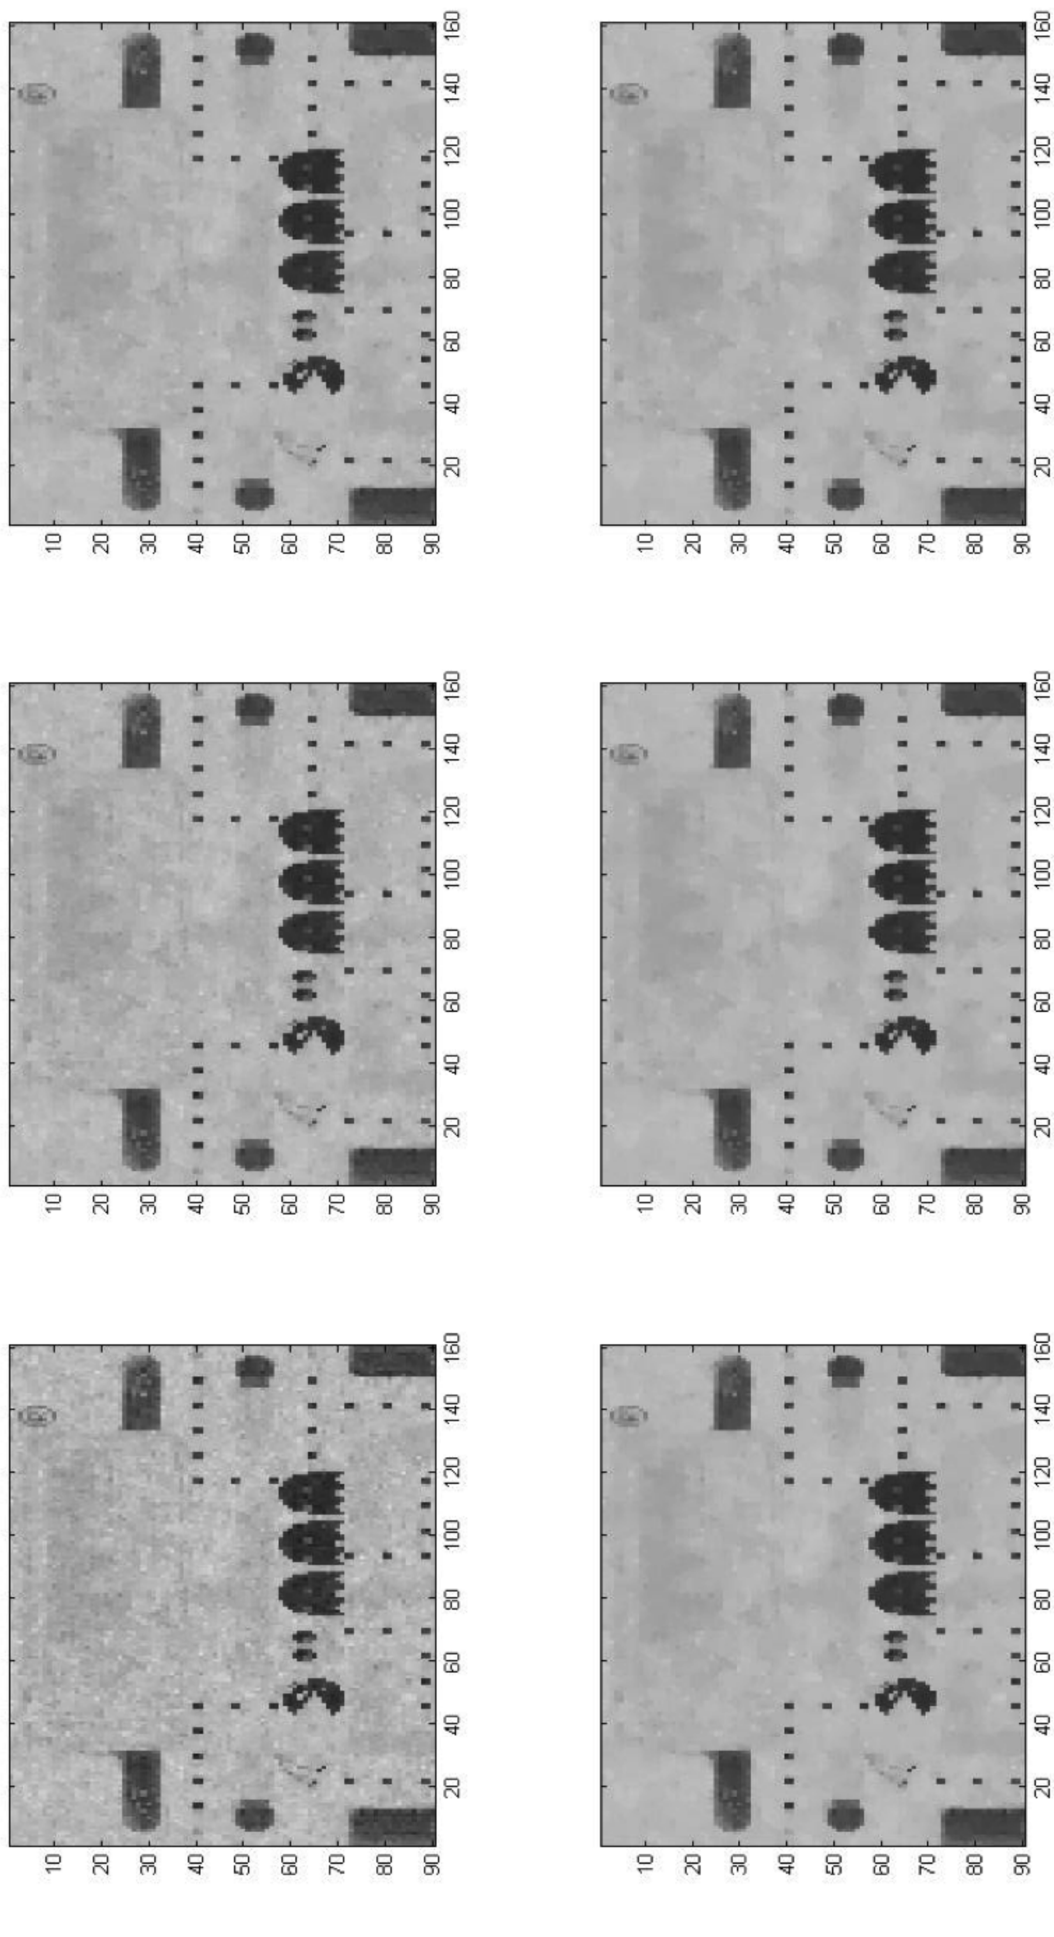

Answer these preliminary questions about the images.

1. In what sense can this be considered a sequence in the metric space of images?

2. Do we have all the terms of the sequence?

3. Does the sequence converge? If so, how can we find the limit of the sequence? Is the limit the original (true) image?

Back to the real line:

To motivate our discussion, consider the following familiar sequence (a_n) of real numbers, where

$$a_n = \sum_{k=1}^{n} \frac{1}{k}.$$

4. Write out the first 4 terms of this **sequence**.

5. What is the difference between successive terms in the sequence? (in other words, for $n \in \mathbb{N}$, what is $a_n - a_{n-1}$?

1. What happens to these differences as $n \to \infty$?

6. Does the sequence (a_n) converge?

Exploration: Examining sequences of images in Matlab

You will now run some denoising code on images and investigate the distances between successive images in Matlab. You will need to put the ZIP file containing this code onto your computer. After unzipping the file, put all files from this code into the matlab/octave directory in which you are working. You can figure out where this is by typing at an octave/matlab prompt the following command.

```
pwd
```

7. Run the code, Lab2Code.m. It will generate a noisy image and denoise that image subject to the default parameters specified.

8. As the denoising code runs you will see a low resolution (in just 2 colors) video of the successive images. After the denoising is complete, the script should produce three figures:

   - Side-by-side comparison of the original image, the noisy image, and the denoised image.
   - A graph of the Euclidean distance between the true image and each image in the denoising sequence.
     (a.) Why do the distances in this graph not go to zero?
   - A graph of the Euclidean distance between successive images in the denoising sequence.
     (b.) What do you notice about the distances in this graph?
     (c.) What does this mean? (And what does it not mean?)

Try changing parameters:

Open the file Lab2Code.m and change only the parameters specified as parameters that you should change. Try running Lab2Code.m with the different parameters. How would you use the results in each of the outputs to try to find the best parameters? You need not actually find the best parameters here. (Be sure to save the results of these trials and include them in your lab write-up.)

Some final thoughts:

In general, we know a sequence converges to a limit if the distance between the limit and the nth term gets arbitrarily small as n approaches infinity. But in our image denoising example, we don't start with a limit image... so how would we tell if our sequence of images is converging? (Note also that in our denoising process, we cannot ever compute the entire infinite sequence, so we never know for sure?) We will come back to these issues when we study the *Cauchy criterion for convergence.*