

ECE 2025 Spring 2000
Lab #5: Digital Images: A/D and D/A

Date: 14–17 Feb 2000

This is *the official* Lab#5 description.

Lab Quiz #2 will be given at the beginning of Lab #6 (next week).

Complete the survey either before or during lab, and have it signed off.

The Warm-up section of each lab must be completed in Lab and the steps marked *Instructor Verification* must also be signed off **during the lab time**. One of the laboratory instructors must verify the appropriate steps by initialing on the **Instructor Verification** line. When you have completed a step that requires verification, simply raise your hand and demonstrate the step to the instructor.

The lab report for this week will an **Informal Lab Report**. Staple the **Instructor Verification** sheet to the end of your lab report as evidence that the appropriate steps were witnessed by the instructor.

The report will **due during the week of 21–24 Feb. at the start of your lab**.

1 Overview

The objective in this lab is to introduce digital images as a second useful signal type. We will show how the A-to-D sampling and the D-to-A reconstruction processes are carried out for digital images. In particular, we will show a commonly used method of image zooming (reconstruction) that gives “poor” results—a later lab will revisit this issue and do a better job.

1.1 Digital Images

In this lab we introduce digital images as a signal type for studying the effect of sampling, aliasing and reconstruction. An image can be represented as a function $x(t_1, t_2)$ of two continuous variables representing the horizontal (t_2) and vertical (t_1) coordinates of a point in space.¹ For monochrome images, the $x(t_1, t_2)$ would be a scalar function of the two spatial variables, but for color images the function $x(\cdot, \cdot)$ would have to be a vector-valued function of the two variables.² Moving images (such as TV) would add a time variable to the two spatial variables.

Monochrome images are displayed using black and white and shades of gray, so they are called *gray-scale* images. In this lab we will consider only sampled gray-scale still images. A sampled gray-scale still image would be represented as a two-dimensional array of numbers of the form

$$x[m, n] = x(mT_1, nT_2) \quad 1 \leq m \leq M, \text{ and } 1 \leq n \leq N$$

where T_1 and T_2 are the sample spacings in the horizontal and vertical directions. Typical values of M and N are 256 or 512; e.g., a 512×512 image which has nearly the same resolution as a standard TV image. In MATLAB we can represent an image as a matrix consisting of M rows and N columns. The matrix entry at (m, n) is the sample value $x[m, n]$ —called a *pixel* (short for picture element).

¹The variables t_1 and t_2 are confusing since they *do not denote time*.

²For example, an RGB color system needs three values for red, green and blue at each spatial location.

An important property of light images such as photographs and TV pictures is that their values are always non-negative and finite in magnitude; i.e.,

$$0 \leq x[m, n] \leq X_{\max}$$

This is because light images are formed by measuring the intensity of reflected or emitted light which must always be a positive finite quantity. When stored in a computer or displayed on a monitor, the values of $x[m, n]$ have to be scaled relative to the maximum value X_{\max} . Usually an eight-bit integer representation is used. With 8-bit integers, the maximum value (in the computer) would be $X_{\max} = 2^8 - 1 = 255$, and there are $2^8 = 256$ gray levels for the display, from 0 to 255.

1.2 Displaying Images

As you will discover, the correct display of an image on a gray-scale monitor can be tricky, especially after some processing has been performed on the image. We have provided the function `show_img.m` in the *DSP First Toolbox* to handle most of these problems,³ but it will be helpful if the following points are noted:



1. All image values must be non-negative for the purposes of display. Filtering may introduce negative values, especially if differencing is used (e.g., a high-pass filter).
2. The default format for most gray-scale displays is eight bits, so the pixel values $x[m, n]$ in the image must be converted to integers in the range $0 \leq x[m, n] \leq 255 = 2^8 - 1$.
3. The actual display on the monitor is created with the `show_img` function.⁴ The `show_img` function will handle the color map and the “true” size of the image. The appearance of the image can be altered by running the pixel values through a “color map.” In our case, we want grayscale so all three primary colors (red, green and blue, or RGB) are used equally, creating what is called a “gray map.” In MATLAB the gray color map is set up via

```
colormap(gray(256))
```

which gives a 256×3 matrix where all 3 columns are equal. The function `colormap(gray(256))` creates a linear mapping, so that each input pixel amplitude is rendered with a screen intensity proportional to its value (assuming the monitor is calibrated). For our experiments, non-linear color mappings would introduce an extra level of complication, so we won't use them.

4. When the image values lie outside the range $[0,255]$, or when the image is scaled so that it only occupies a small portion of the range $[0,255]$, the display may have poor quality. In this lab, we will use `show_img.m` to *automatically rescale the image*: This requires a linear mapping of the pixel values:⁵

$$x_s[m, n] = \mu x[m, n] + \beta$$

The scaling constants μ and β can be derived from the min and max values of the image, so that all pixel values are recomputed via:

$$x_s[m, n] = \left\lfloor 255.999 \left(\frac{x[m, n] - x_{\min}}{x_{\max} - x_{\min}} \right) \right\rfloor$$

where $\lfloor x \rfloor$ is the floor function, i.e., the greatest integer less than or equal to x .

³If you have the image processing toolbox, then the function `imshow.m` can be used.

⁴If the MATLAB function `imagesc.m` is used to display the image, two features will be missing: (1) the color map may be incorrect because it will not default to gray, and (2) the size of the image will not be a true pixel-for-pixel rendition of the image on the computer screen.

⁵The MATLAB function `show_img` has an option to perform this scaling while making the image display.

1.3 Printing Multiple Images on One Page

The phrase “what you see is what you get” can be elusive when dealing with images. It is **VERY TRICKY** to print images so that the hard copy matches exactly what is on the screen, because there is usually some interpolation being done by the printer or program that is handling the images. One way to think about this in signal processing terms is that the screen is one kind of D-to-A and the printer is another kind, but they use different (D-to-A) reconstruction methods to get the continuous-time (analog) output image that you see.

Furthermore, if you try to put two images of different sizes into subplots of the same MATLAB figure, it won't work. Therefore, you should display your images in separate figure windows. In order to get a printout with **MULTIPLE IMAGES ON THE SAME PAGE**, use the following procedure:

1. In MATLAB, use `show_img` and `trusize` to put your images into separate figure windows at the correct pixel resolution.
2. Use the Windows-95 program called PAINT to assemble the different images onto one page. This program can be found under Accessories.
3. For each MATLAB figure window, do ALT-PRINT-SCREEN which will copy the active window contents to the clipboard.
4. After each “window capture” in step 3, paste the clipboard contents into PAINT.
5. Arrange the images so that you can make a comparison for your lab report.
6. Print the assembled images from PAINT to a printer.

2 Warm-up

The instructor verification sheet may be found at the end of this lab.

2.1 Warm-up: Display of Images

You can load the images needed for this lab from *.mat files. Any file with the extension *.mat is in MATLAB format and can be loaded via the `load` command. To find some of these files, look for *.mat in the *DSP First* toolbox or in the MATLAB directory called `toolbox/matlab/demos`. Some of the image files are named `lenna.mat`, `echart.mat` and `zone.mat`, but there are others within MATLAB's demos. The default size is 256×256 , but alternate versions are available as 512×512 images under names such as `lenna512.mat` and `zone512.mat`. After loading, use the command `whos` to determine the name of the variable that holds the image and its size.

Although MATLAB has several functions for displaying images on the CRT of the computer, we have written a special function `show_img()` for this lab. It is the visual equivalent of `soundsc()`, which we used when listening to speech and tones; i.e., `show_img()` is the “D-to-C” converter for images. This function handles the scaling of the image values and allows you to open up multiple image display windows. Here is the help on `show_img`:

```
function [ph] = show_img(img, figno, scaled, map)
%SHOW_IMG    display an image with possible scaling
% usage:    ph = show_img(img, figno, scaled, map)
%          img = input image
%          figno = figure number to use for the plot
%                  if 0, re-use the same figure
%                  if omitted a new figure will be opened
```



CD-ROM

IMAGE
DATA
FILES



CD-ROM

show_im.m

```

% optional args:
%   scaled = 1 (TRUE) to do auto-scale (DEFAULT)
%           not equal to 1 (FALSE) to inhibit scaling
%   map = user-specified color map
%   ph = figure handle returned to caller
%-----

```

Notice that unless the input parameter `figno` is specified, a new figure window will be opened.

2.2 Display Test

In order to probe your understanding of image display, do the following simple displays:

- (a) Generate a simple test image in which all of the rows are identical by using the following outer product:

```
xpix = ones(256,1)*cos(2*pi/8*(1:256));
```

Display the image and explain the gray-scale pattern that you see. How wide are the bands in number of pixels? How can you predict that width from the formula for `xpix`?

Instructor Verification (separate page)

- (b) Now load and display the 256×256 “lenna” image from `lenna.mat`. The command `load lenna` will put the sampled image into the array `xx`. Use `whos` to check the size of `xx` after loading. When you display the image it might be necessary to set the colormap via `colormap(gray(256))`.
- (c) Use the colon operator to extract the 200th row of the “lenna” image, and make a plot of the 200th row as a 1-D discrete-time signal. Observe that the range of signal values is between 0 and 255. Which values represent white and which ones black?

Instructor Verification (separate page)

3 Lab Exercises: Sampling, Aliasing and Reconstruction

3.1 Sampling of Images

Images that are stored in digital form on a computer have to be sampled images because they are stored in an $M \times N$ array (i.e., a matrix). The sampling rate in the two spatial dimensions was chosen at the time the image was digitized (in units of samples per inch if the original was a photograph). For example, the image might have been “sampled” by a scanner where the resolution was chosen to be 300 dpi (dots per inch).⁶ If we want a different sampling rate, we can simulate a *lower* sampling rate by simply throwing away samples in a periodic way. For example, if every other sample is removed, the sampling rate will be halved (in our example, the 300 dpi image would become a 150 dpi image). Usually this is called *sub-sampling* or *down-sampling*.

Down-sampling: If the vector `x1` represents a signal such as a row of an image, we can reduce the sampling rate by a factor of 4 by simply taking every 4th sample. In MATLAB this is easy with the colon operator, i.e., `xdown = x1(1:4:length(x1))`. The vector `xdown` will be one fourth the length of `x1` (use `whos` to verify this fact).

⁶The Sampling Theorem applies to digital images, so there is a *Nyquist Rate* that depends on the maximum *spatial* frequency in the image.

- (a) *Down-sampling* throws away samples, so it will shrink the size of the image. This is what is done by the following scheme:

$$xp = xx(1:p:M, 1:p:N) ;$$

when we are downsampling by a factor of p . One potential problem with down-sampling is that aliasing might occur. This can be illustrated in a dramatic fashion with the `barbara` image, and to a lesser degree with the `lenna` image.

Load the `barbara.mat` file which has the image stored in a variable called `barb`. When you check the size of the image, you'll find that it is not square. Now down-sample the `barbara` image by a factor of 3. What is the size of the down-sampled image? Notice the aliasing in the down-sampled image, which is surprising since no new values are being created by the down-sampling process. Describe how the aliasing appears visually.⁷

- (b) This part is challenging: explain why the aliasing happens in the `barbara` image by using a “frequency domain” explanation. In other words, estimate the frequency of the features that are being aliased. Give this frequency will be in cycles per pixel. Can you relate your frequency estimate to the Sampling Theorem?

You might try zooming in on a very small region of both the original and downsampled images.

3.2 Reconstruction of Images

When an image has been sampled, we can fill in the missing samples by doing interpolation. For images, this would be analogous to the examples shown in Chapter 4 for sine-wave interpolation which is part of the reconstruction process in a D-to-A converter. We could use a “square pulse” or a “triangular pulse” or other pulse shapes for the reconstruction.

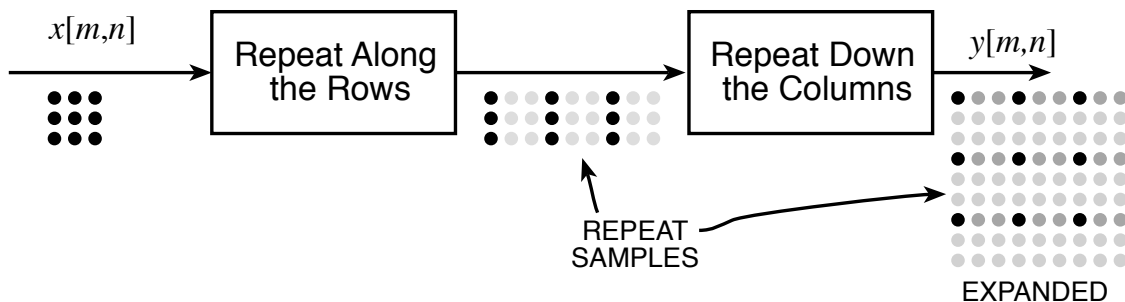


Figure 1: 2-D Interpolation broken down into row and column operations: the gray dots indicate repeated data values created by a zero-order hold; or, in the case of linear interpolation, they are the interpolated values.

For these reconstruction experiments, use the `barbara` image, down-sampled by 3 (Section 3.1). You will have to generate this by loading in the image from `barbara.mat` to get the image which is in the array called `barb`. The down-sampled `barbara` image should be created and stored in the variable `xx3`. The objective will be to reconstruct an approximation to the original `barbara` image, which is 484×393 , from the smaller down-sampled image.

⁷One difficulty with showing aliasing is that we must display the pixels of the image exactly. This almost never happens because most monitors and printers will perform some sort of interpolation to adjust the size of the image to match the resolution of the device. In MATLAB we can override these size changes by using the function `truesize` which is part of the Image Processing Toolbox. In the *DSP First Toolbox*, an equivalent function called `trusize.m` is provided.

- (a) The simplest interpolation would be reconstruction with a square pulse which produces a “zero-order hold.” Here is a method that works for a one-dimensional signal (i.e., one row or one column of the image), assuming that we start with a row vector `xr1`, and the result is the row vector `zz`.

```
xr1 = (-2).^ (0:6);
L = length(xr1);
nn = ceil((0.999:1:3*L)/3);    %<-- Round up to the integer part
zz = xr1(nn);
```

Plot the vector `zz` to verify that it is a zero-order hold version derived from `xr1`. Explain what values are contained in the indexing vector `nn`. Your lab report should include the explanation for this part, but no plots are required.

- (b) Now return to the down-sampled `barbara` image, and process all the rows of `xx3` to fill in the missing points. Use the zero-order hold idea from part (a). Call the result `xholdrows`. Display `xholdrows` as an image, and compare it to the downsampled image `xx3`; compare the size of the images as well as their content.
- (c) Now process all the columns of `xholdrows` to fill in the missing points in each column and call the result `xhold`. Compare the result (`xhold`) to the original image `barb`. Include your code for parts (b) and (c) in the lab report.
- (d) *Linear interpolation* can be done in MATLAB using the `interp1` function (that’s “interp-one”).

When unsure about a command, use `help`.

Its default mode is linear interpolation, which is equivalent to using the ‘`*linear`’ option, but `interp1` can also do other types of polynomial interpolation. Here is an example on a 1-D signal:

```
n1 = 0:6;
xr1 = (-2).^n1;
tti = 0:0.1:6;    %-- locations between the n1 indices
xr1lin = interp1(n1,xr1,tti);
stem(tti,xr1lin)
```

Carry out a linear interpolation operation on both the rows and columns of the image. This requires two calls to the `interp1` function, because one call will only process all the columns of a matrix.⁸ Name the interpolated output image `xlin`. Include your code for this part in the lab report.

- (e) Compare `xlin` to the original image `barb`. Comment on the visual appearance of the “reconstructed” image versus the original; point out differences and similarities. Can the reconstruction (i.e., zooming) process remove the aliasing effects from the down-sampled `barbara` image?
- (f) Compare the quality of the linear interpolation result to the zero-order hold result. Point out regions where they differ and try to justify this difference by estimating the local frequency content. In other words, look for regions of “low-frequency” content and “high-frequency” content and see how the interpolation quality is dependent on this factor.

A couple of questions to think about: Are edges low frequency or high frequency features? Are the stripes low frequency or high frequency features?

Comment: You might use MATLAB’s zooming feature to show details in a small patches of the output image. However, be careful because zooming does its own interpolation, probably a zero-order hold.

⁸Use a matrix transpose in between the interpolation calls. The transpose will turn rows into columns.

3.3 More about Images in MATLAB (Optional)

This section is included for those students who might want to relate these MATLAB operations to previous experience with software such as *Photoshop*. There are many image processing functions in MATLAB. For example, try the help command:

```
help images
```

for more information, but keep in mind that the Image Processing Toolbox, which is available in the CoC labs, may not be on your computer.

3.3.1 Zooming in Software

If you have used an image editing program such as Adobe's *Photoshop*, you might have observed how well or how poorly image zooming (i.e., interpolation) is done. For example, if you try to blow up a JPEG file that you've downloaded from the web, the result is usually disappointing. Since MATLAB has the capability to read lots of different formats, you can apply the image zooming via interpolation to any photograph that you can acquire. The MATLAB function for reading JPEG images is `imread()` which would be invoked as follows:

```
xx = imread('foo.jpg','jpeg');
```

Since `imread()` is part of the image processing toolbox, this test can be done in the CoC computer labs, but may not be possible on your home computer.

3.3.2 Warnings

Images obtained from JPEG files might come in many different formats. Two precautions are necessary:

1. If MATLAB loads the image and stores it as 8-bit integers, then MATLAB will use an internal data type called `uint8`. The function `show_img()` cannot handle this format, but there is a conversion function called `double()` that will convert the 8-bit integers to double-precision floating-point for use with filtering and processing programs.

```
yy = double(xx);
```

You can convert back to 8-bit values with the function `uint8()`.

2. If the image is a color photograph, then it is actually composed of three "image planes" and MATLAB will store it as a 3-D array. For example, the result of `whos` for a 545×668 color image would give:

Name	Size	Bytes	Class
xx	545x668x3	1092180	uint8 array

In this case, you should use MATLAB's image display functions such as `imshow()` to see the color image. Or you can convert the color image to gray-scale with the function `rgb2gray()`. For more information on the image processing functions in MATLAB, try help:

```
help images
```

Lab #5
ECE-2025
Spring-2000
INSTRUCTOR VERIFICATION PAGE

Staple this page to the end of your Lab Report.

Name: _____ Date of Lab: _____

Complete the on-line survey in Web-CT:

Verified: _____ Date/Time: _____

Part 2.2(a) Generate and display a digital image. Explain the width of bands in the test image formed from the “outer product” of ones with a cosine. Give the width of the bands in number of pixels and explain how you can predict that width from the formula for x_{pix} ?

Verified: _____ Date/Time: _____

Part 2.2(c) Load and display a digital image; extract one row of the image.

Verified: _____ Date/Time: _____