AMATH 231 Supplementary Notes

Representation of images by scalar- and vector-valued functions

`These notes have been taken from the instructor's lecture notes (Lecture 1) in AMATH 391 Fall 2017`

Black-and-white images as scalar-valued functions

In this section, we review the idea that images may be represented by real-valued functions of two real variables. (This section was presented in Lecture 1 of this AMATH 231 course.)

Images may be considered to be two-dimensional signals. An "ideal image," as approximated by a photograph, may be represented mathematically by a function of two spatial variables, i.e., $f(x, y)$, where x and y are continuous variables over a finite set $D \subset \mathbb{R}^2$, the domain of the function, i.e., $(x, y) \in D$. For a **black-and-white** image, $f(x, y)$ assumes a real and typically non-negative value – the so-called **greyscale value** – that characterizes the "greyness" at a point (x, y) of the image. Mathematically, f is a real-valued function, i.e., $f : \mathbb{R}^2 \to \mathbb{R}$.

For simplicity, let us assume that the domain of f is given by $x, y \in [0, 1]$, which we may also write as $(x, y) \in [0, 1]^2$. As well, assume that the range of f is the interval $[0, 1]$, i.e., $f : [0, 1] \to [0, 1]^2$. Then the value 0 will represent black and the value 1 will represent white. An intermediate value, i.e., $0 < f(x, y) < 1$, will represent some shaded grey value. The graph of the image function, $z = f(x, y)$, then may be viewed as a representation of the image, as shown in the figure below on the right.

Colour images as vector-valued functions

An ideal **colour image** is represented mathematically by a **vector-valued function**. At each point $(x, y) \in [0, 1]$, are defined three colour values, namely, **red**, **green** and **blue**. The combination of these three **primary colours** produces the colour associated with (x, y). Mathematically, f is a mapping from $\mathbb{R}^2$ to $\mathbb{R}^3$ having the form,

$$f(x, y) = (r(x, y), g(x, y), b(x, y)), \tag{1}$$

where $r(x, y)$, $g(x, y)$ and $b(x, y)$ denote, respectively, the red, green and blue values at (x, y).

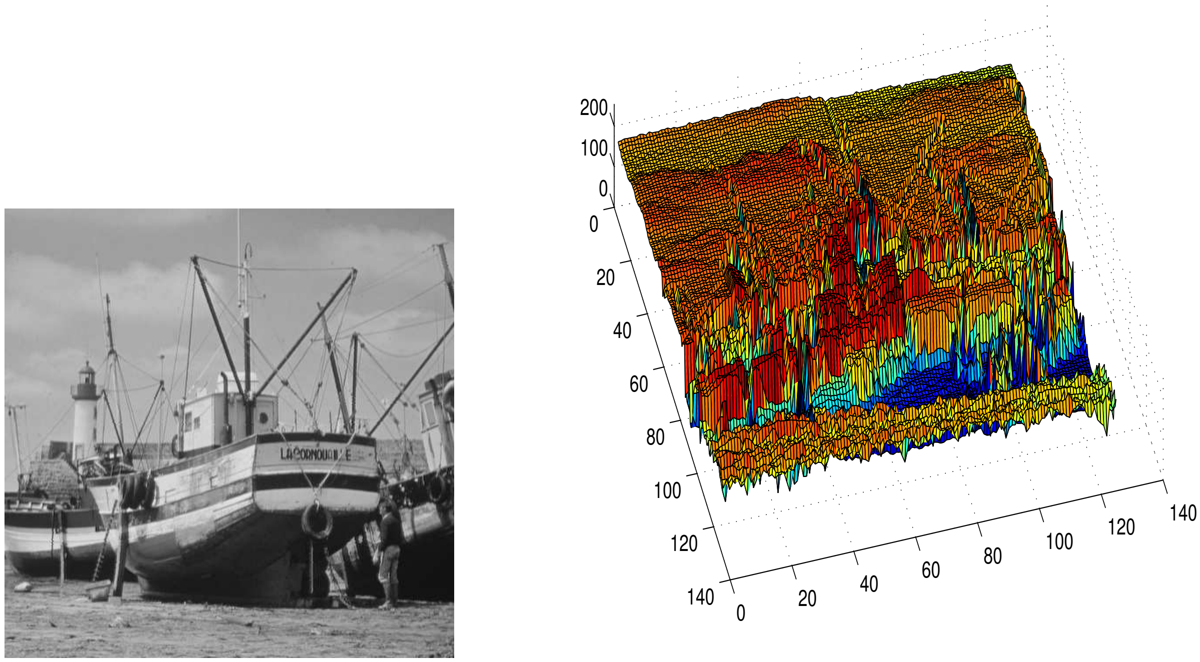

Left: The standard test-image, *Boat*, a 512×512-pixel digital image, 8 bits per pixel. Each pixel assumes one of 256 greyscale values between 0 and 255. **Right:** The *Boat* image, viewed as an image function $z = f(x, y)$. The red-blue spectrum of colours is used to characterize function values: Higher values are more red, lower values are more blue.

Digital images

Digital images are two-dimensional arrays that represent samplings of the image function $f(x, y)$. **Black-and-white digital** images are represented by $n_1 \times n_2$ matrices, $\mathbf{u} = \{u_{ij}\}$. (As the caption indicates, the *Boat* image of the previous figure is a black-and-white digital image.) The entry u_{ij} of this matrix – usually written as $u[i, j]$ in image processing literature – represents the greyscale values at the (i, j) pixel, $1 \leq i \leq n_1$, $1 \leq j \leq n_2$. The greyscale values of digital images also assume discrete values so that they may easily be stored in digital memory. The typical practice is to allocate n bits of memory for each greyscale value so that a total of 2^n values, namely, $\{0, 1, 2, \cdots, 2^n - 1\}$ are employed. In most applications, $n = 8$, i.e., 8-bit images, implying a set of 256 greyscale values ranging from 0 to 255. This is found to be more than sufficient for the human visual system.

Colour digital images will be represented by three matrices, $\mathbf{r} = \{r[i, j]\}$, $\mathbf{g} = \{g[i, j]\}$ and $\mathbf{b} = \{b[i, j]\}$, which represent, respectively, the red, green and blue values at the (i, j) pixel. As you may recall from an earlier Physics or Science courses, any colour can be generated by means of an appropriate combination of these three **primary colours**.

In the top left of the figure below is presented the digital colour image, *Sailboat-lake*, a 512×512-pixel image, 24 bits per pixel, 8 bits per colour. At each pixel, 8 bits, i.e., 256 values ranging from 0 to 255, are used to store the intensity of each of the three colours – red, green and blue. The red, green and blue component images are shown in the figure. Note that they are displayed as black-and-white images, with $0 =$ black and $255 =$ white.

A closer look at the component images will show why particular regions of particular primary colours are either low or high in magnitude. One would expect that the blue components for pixels representing the blue sky in the colour image would have a higher intensity than red and green components. This, in turn, would imply that the black-and-white image representing the blue component would be lighter/whiter in the blue sky regions, which is observed to be the case. One can also draw similar conclusions for the red components in reddish regions of the colour image as well as green components in greener regions of the colour image.

Hyperspectral images

The red, green and blue values of an image at a point/pixel may be viewed as the reflectance values for a particular **sampling** of the visual electromagnetic spectrum at three wavelengths, $\lambda_r > \lambda_g > \lambda_b$. As mentioned earlier, three colours are sufficient since any visible colour may be generated from a combination of these three primary colours.

That being said, in other applications, e.g., remote sensing, a much greater sampling of the electromagnetic spectrum is performed. For example, the Airborne Visible/Infrared Integrating Spectrometer (AVIRIS), performs a sampling of 224 wavelengths in the visible and infrared regions of the electromagnetic spectrum. Such high sampling is performed in order to determine the composition or nature of regions being photographed. Various soils, for example, based upon their mineral composition exhibit different reflectance spectra, i.e., the "shape" of the 224-vectors. The same may be said for different types of vegetation, etc.. Some years back, when satellite imagery was in its infancy, and a much lower degree of sampling was performed, such images were known as **multispectral images**. Now, with much greater degrees of sampling, these images are known as **hyperspectral images**. Once again, a hyperspectral image may be viewed as a vector-valued function: At each pixel location $[i, j]$, which represents a particular region of the earth's surface, the hyperspectral image function f

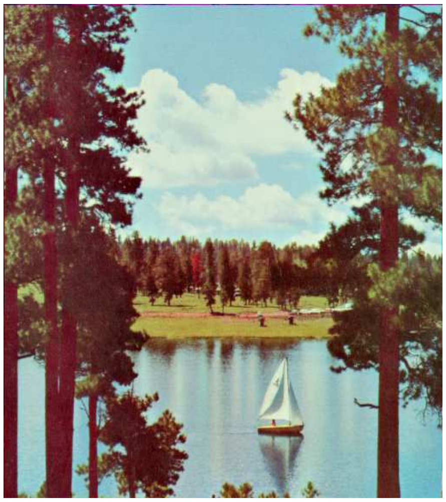

RGB colour image

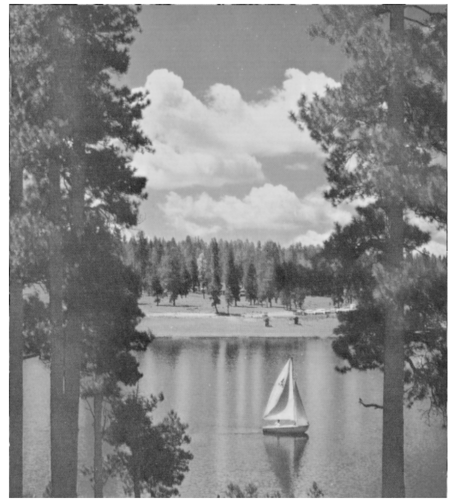

Red component

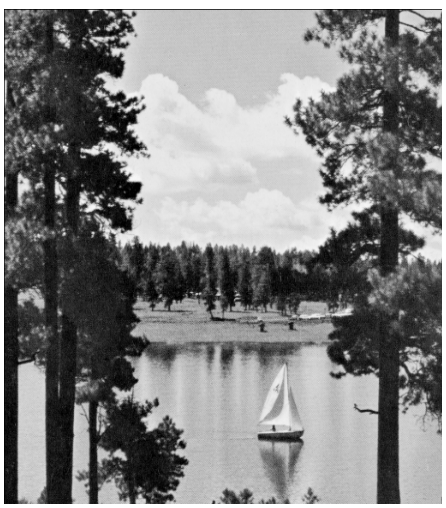

Green component

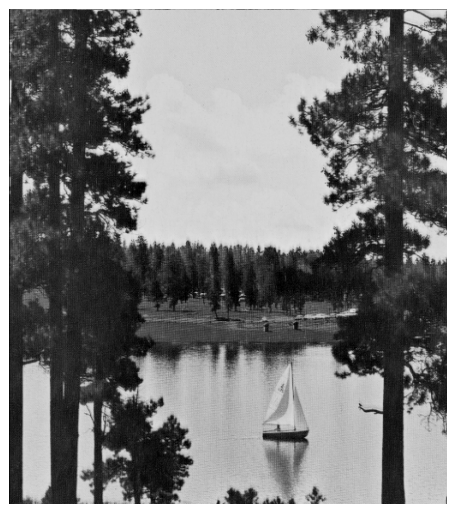

Blue component

The standard digital colour test-image, *Sailboat-lake*, 512×512-pixel, 24 bits per pixel (8 bits per colour), along with its red, green and blue component images.

is an M-vector, i.e.,

$$f[i,j] = (f_1[i,j], f_2[i,j], \cdots, f_M[i,j]) \,. \tag{2}$$

This M-vector defines the **spectral function** of the region represented by the pixel $[i,j]$. As mentioned earlier, spectral functions can contain a great deal of information about the chemical composition of regions.

The component functions $f_k(x, y)$ corresponding to different wavelengths are often referred to as **channels**. Each channel represents an image of the particular region on the earth taken at a particular wavelength.

A pictorial representation of the "stacking up" of many channels to form a hyperspectral **"data cube"** is shown in the figure below.

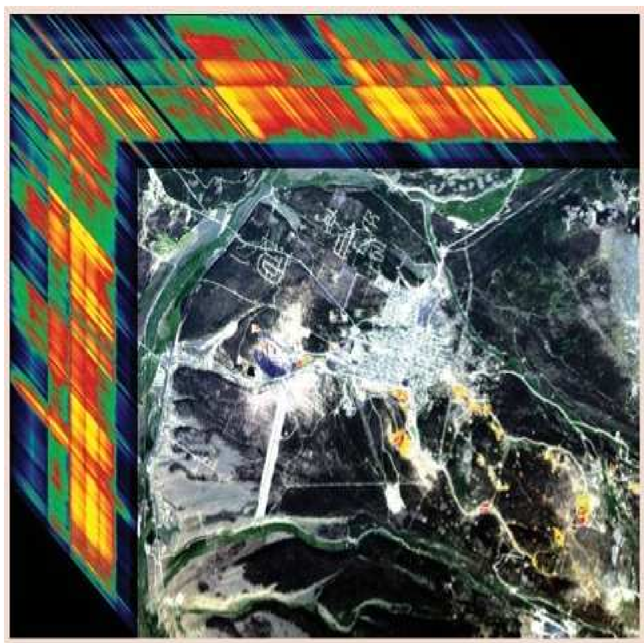

A pictorial representation of the "stacking up" of images – or channels – corresponding to different wavelengths to form a hyperspectral image.

Another type of hyperspectral imaging: Diffusion magnetic resonance imaging (dMRI)

You have most probably heard of **magnetic resonance imaging** (MRI), which is based on the so-called **magnetic moment** of the hydrogen atom nucleus, namely, the proton. Very briefly, a proton will interact with an external magnetic field due to its intrinsic **magnetic moment**. As you know, water, "H_2O", which is composed of hydrogen and oxygen atoms, is present virtually everywhere in living tissue. Different regions in a human body, e.g., different organs, tissues, etc., represent different structural and biochemical environments for the water molecules within those regions. As such, protons

The nature of spectral functions

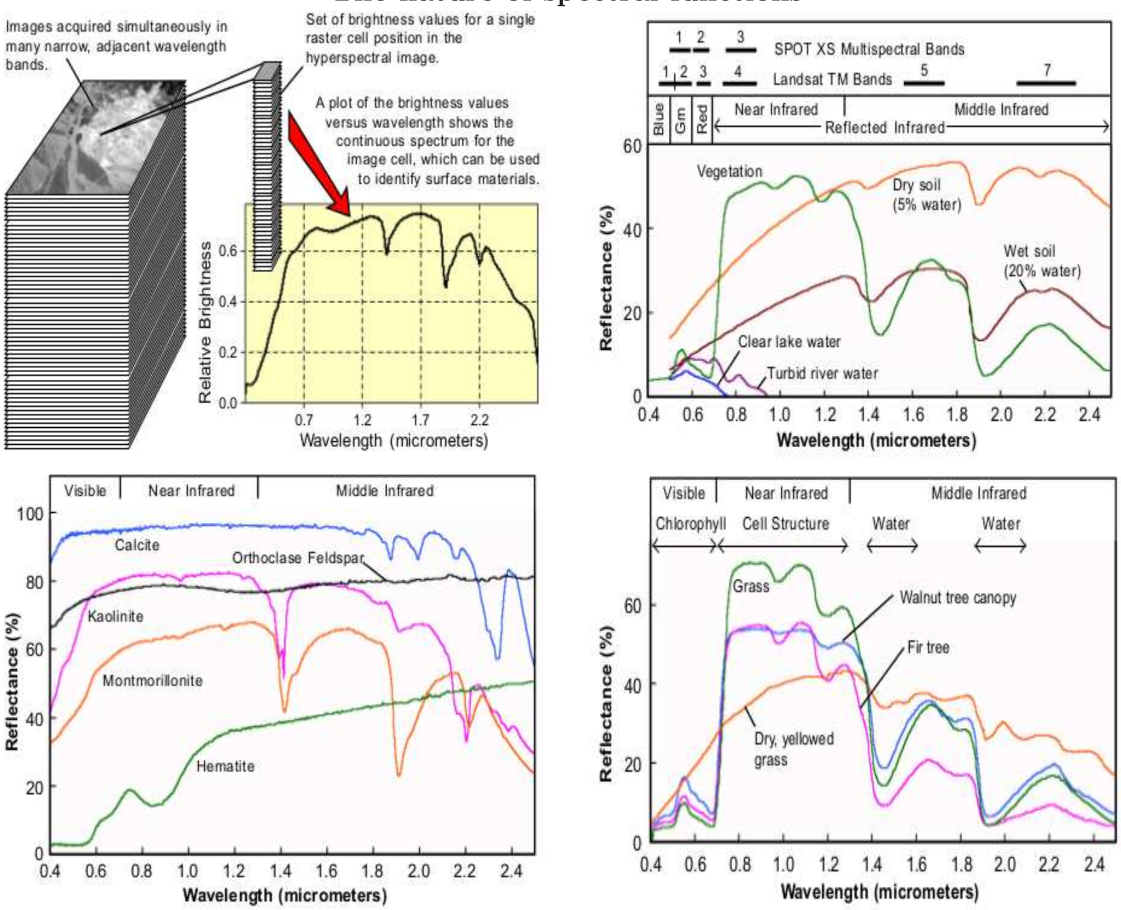

in water molecules from different regions will respond differently to an external magnetic field. (A little more precisely – protons from different regions will have different rates of **spin relaxation** in the presence of a constant external magnetic field.) In a very clever manner, magnetic resonance imaging uses these differences in response (relaxation) to produce a two- or three-dimensional pictorial representation of the interior of a human body (or whatever is being imaged).

The more recently developed technique of **diffusion magnetic resonance imaging** (DMRI) also exploits the magnetic moment of protons. DMRI is able to detect the motion of collections of water molecules in local regions of the body under observation. Realistically, because of limitations in resolution, the characterization of the motion is limited to a finite number of directions. The net result of this procedure is that at a 3D pixel location (i, j, k) in the interior of the body being observed, one can estimate the probability that water molecules at (i, j, k) will move (diffuse) in each of M directions. Once again, the result is a vector-valued image function,

$$u[i, j, k] = (u_1[i, j, k], u_2[i, j, k], \cdots, u_M[i, j, k])\,. \tag{3}$$

This is illustrated in the figure below.

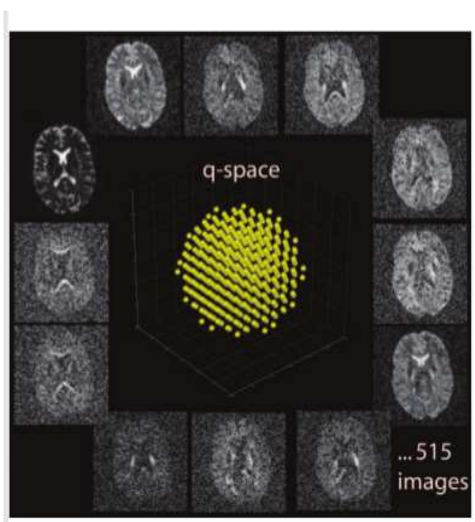

Figure 12. Series of diffusion-weighted MR brain images obtained with variations in the direction and strength of the diffusion gradient in the pulsed gradient SE sequence. Each image shows the signal sampled at one point in q-space (one yellow dot). Every sampling point in q-space corresponds to a specific direction and strength of the diffusion gradient.

From *Understanding Diffusion MR Imaging Techniques: From Scalar Diffusion-weighted imaging to Diffusion Tensor Imaging and Beyond*, by P. Hagmann *et al.*, Radiographics 2006, 26:S205-S223. Published online `10.1148/rg.26si065510`.

One fascinating application of DMRI is in neurobiology – the ability to produce maps of neural connections in the brain. It seems natural that there will be a greater probability for water molecules inside a neuron to travel in the tubular direction of the neuron as opposed to through its boundaries. At each pixel, one can assign a most probable direction for water molecules to exit the region represented by the pixel. The result is a vector field defined over all pixels. One can now "connect the arrows" – essentially the procedure of finding field lines of the vector field examined in this course – to obtain connectivity maps such as the one in the figure below. These connectivity maps are called **connectomes**.

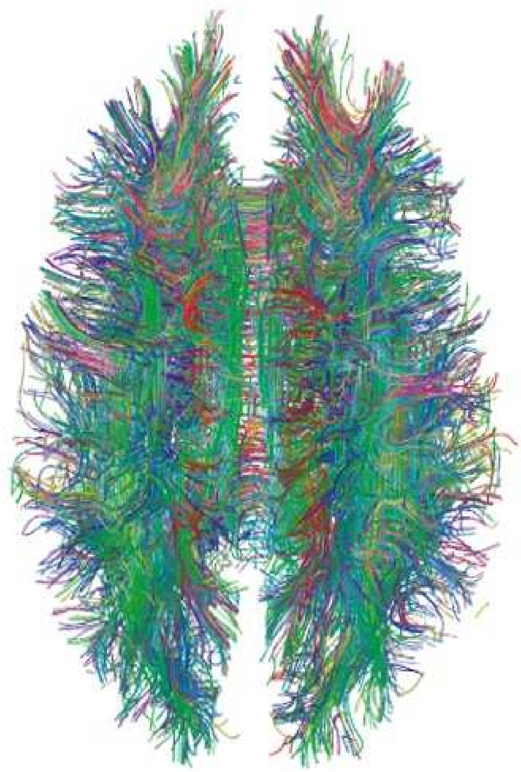

A typical "connectome," a pictorial representation of the connectivity of neurons in the human brain.