Menu September 16, 2009

Topics:

Image Characterization
— continuous case, $i(x, y)$
— discrete case, $I(X, Y)$
When is $I(X, Y)$ an *exact* characterization of $i(x, y)$?

Reading:

Text: Forsyth & Ponce Chapter 7
Papers: 2.1, 2.2, 2.4, 2.5

Reminders:

Ian Mitchell's Matlab tutorial is **today**, 5pm, DMP 101
Web site: `http://people.cs.ubc.ca/~woodham/cpsc505/`
Newsgroup: ubc.courses.cpsc.505

Today's "Fun" Example: Automimicry

The Peablue (*Lampides boeticus*) is a small butterfly. Notice the eyespot and pseudo-antennae on the hindwing

Image credit:

`http://en.wikipedia.org/wiki/Image:Peablue_October_2007_Osaka_Japan.jpg`

The image in today's example was selected Wikipedia "Picture of the day," Saturday, March 8, 2008. Visit Wikipedia at the URL given above to view the full resolution (3,320 × 2,222 pixel, 4.09MB) version available there.

Here is a brief summary of the take–away points from Lecture 2.

Lecture 2: Re-cap

- We are concerned with the relationships among 3D worlds, 2D images and perception
- Two principal motivations:
  — understanding vision
  — making machines see
- We take a "physics-based" approach
- We are careful to distinguish between the 2D image domain and the 3D scene domain
- In CPSC 505, we distinguish terminology according to the:
  — *problem addressed*
  — *tools used*
- We prefer the term *computational vision*

Despite the fact that the term "computational vision" doesn't appear in the title of any of the courses, CPSC 425, 505 and 525, that is the term we prefer. The problem we are interested in is "vision," whether by biological system or by machine. The approach we take is "computational." Thus, the term "computational vision" identifies both the problem of interest and the tools used. For the same reason, we prefer the term "computational intelligence" to "artificial intelligence." Note: Our research group is the "Laboratory for Computational Intelligence." Earlier, when the focus was more narrow, we were the "Laboratory for Computational Vision."

Now for today's topic. . .

We ask the question, "What is an image?" The answer we give today is mathematical. (ASIDE: We return to "physics" later).

What is an Image?

- Today's lecture provides a *mathematical characterization*

  — to understand how to represent images digitally

  — to understand how to compute with images

First, we consider the continuous case. This both corresponds to our intuition and simplifies the development of a mathematical theory of image processing.

Continuous Case

- The term "image" suggests a 2D surface whose *appearance* varies from point–to–point

  — the 2D surface typically is a plane (but might be curved)
- Appearance can be B&W or colour
- In B&W, variation in appearance can be described by a single parameter corresponding to the amount of light reaching the image at a given point in a given time

  — this is what, for example, B&W film records

Continuous Case (cont'd)

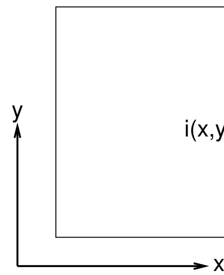

Denote a B&W image as a function, $i(x, y)$, where x and y are *spatial variables*

ASIDE: The convention for today's lecture is to use lower case letters for the continuous case and upper case letters for the discrete case

Continuous Case (cont'd)

Observations:

- $i(x, y)$ is a *real-valued function* of *real spatial variables*, x and y
- $i(x, y)$ is *bounded above and below*. That is,

$$0 \ \leq \ i(x, y) \ \leq \ M$$

  for some maximum brightness, M
- $i(x, y)$ is *bounded in spatial extent*. That is, $i(x, y)$ is non-zero (i.e., strictly positive) over, at most, a bounded region

Our intuition is that an image function, $i(x, y)$, need only be defined over a finite domain. The notion of "bounded spatial extent" captures this intuition. Mathematically, we need to allow both spatial variables, x and y, to be unbounded. That is, $-\infty \ \leq \ x \ \leq \ \infty$ and $-\infty \ \leq \ y \ \leq \ \infty$. To handle this, we embed the finite domain of interest into an infinite background, with $i(x, y) \ = \ 0$ in the background. Conceptually, as was done to create the slide, we draw a (rectangular) box such that $i(x, y) \ = \ 0$ outside the box. Note: Of course, there still can be locations within the box at which $i(x, y) = 0$. The constraint within the box is that $0 \ \leq \ i(x, y) \ \leq \ M$.

Continuous Case (cont'd)

Some additional considerations:

- Images also can be considered a function of time. Then, we write $i(x, y, t)$ where x and y are spatial variables and t is a *temporal variable*
- To make the dependence of brightness on wavelength explicit, we write $i(x, y, t, \lambda)$ where x, y and t are as above and where λ is a *spectral variable*
- More commonly, we think of "colour" already as discrete and write

$$i_R(x, y)$$
$$i_G(x, y)$$
$$i_B(x, y)$$

  for specified colour channels, R, G and B

ASIDE: There are times when we also want to consider $i(x, y)$ to be a function of the state of polarization of the light or, in the case of coherent (laser) optics, a function of the phase of the light. This is beyond the scope of CPSC 505.

Discrete Case

Idea: Superimpose (regular) grid on continuous image

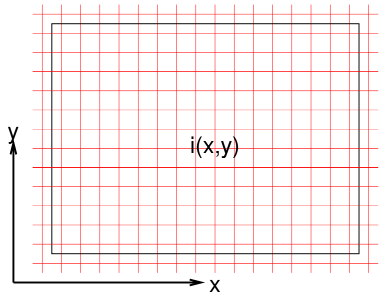

Sample the underlying continuous image according to the *tessellation* imposed by the grid

Discrete Case (cont'd)

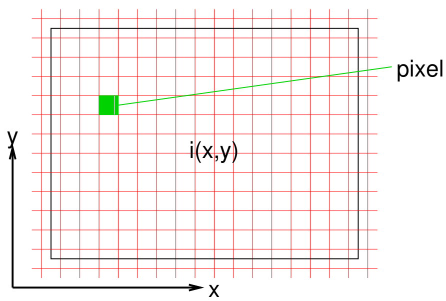

Each grid cell is called a picture element (or *pixel*)

There are a finite number of grid lines in each direction

Without loss of generality, let the grid lines correspond to integer values of x and y

We can refer to the $i\,j^{th}$ pixel in the discrete image

Discrete Case (cont'd)

Denote the discrete image as

$$I(X, Y)$$

$$I(i, j)$$

$$I_{i\,j}$$

Important: We can store the pixels in a matrix or an array

RECALL: The convention for today's lecture is to use lower case letters for the continuous case and upper case letters for the discrete case

Discrete Case (cont'd)

Question: How to sample?

- Sample brightness at point?
- "Average" brightness over entire pixel?

Answer:

- Point sampling is useful for theoretical development
- Area-based sampling occurs in practice

ASIDE: Typically, one can model area-based sampling as the point sampling of a "blurred" version of the original image

Discrete Case (cont'd)

Question: What about the brightness samples themselves?

Answer: We make the values of $I(X, Y)$ discrete, as well

RECALL: $0 \leq i(x, y) \leq M$

We divide the range $[0, M]$ into a finite number of equivalence classes. This is called *quantisation* (American: *quantization*)

The values are called *grey-levels* (American: *gray-levels*)

Discrete Case (cont'd)

Quantisation is a topic in its own right

For now, a simple linear scheme is sufficient

Suppose n bits–per-pixel are available. One can divide the range $[0, M]$ into evenly spaced intervals as follows:

$$i(x, y) \quad \rightarrow \quad \left\lfloor \frac{i(x, y)}{M} (2^n - 1) + 0.5 \right\rfloor$$

where $\lfloor \ \rfloor$ is floor (i.e., greatest integer less than or equal to)

Typically $n = 8$ resulting in grey-levels in the range $[0, 255]$

Sampling Theory (Informal)

Question: When is $I(X, Y)$ an *exact characterization* of $i(x, y)$?

Question (modified): When can we reconstruct $i(x, y)$ *exactly* from $I(X, Y)$?

Intuition: Reconstruction involves some kind of *interpolation*

Heuristic: When in doubt, consider simple cases

Sampling Theory (cont'd)

Case 0: Suppose $i(x, y) = k$ (with k one of our grey-levels)

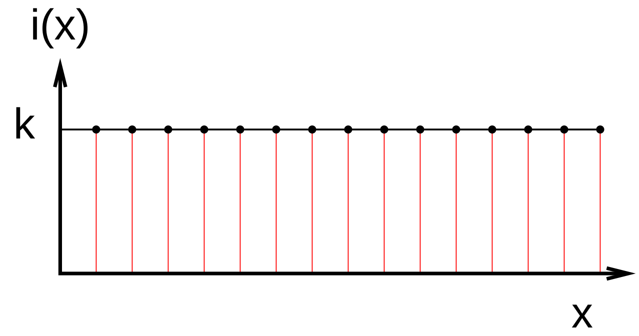

This case is easy!
$I(X, Y) = k$. Any standard interpolation function would give $i(x, y) = k$ for non-integer x and y

Sampling Theory (cont'd)

Case 1: Suppose $i(x, y)$ has a discontinuity not falling precisely at integer x, y

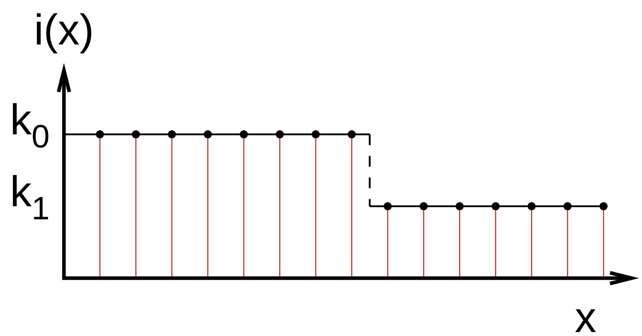

This case is impossible!
We can not reconstruct $i(x, y)$ exactly because we can never know exactly where the disconinuity lies

Sampling Theory (cont'd)

Question: How do we close the gap between "easy" and "impossible?"

Answer: Next, we build intuition based on informal argument

Tell me more!