CS 308 Data Structures

Spring 2003 - Dr. George Bebis

Programming Assignment 1

Due date: 2/25/03

In this assignment, you will write C++ code to read, write, and manipulate images. The objectives are the following:

- Familiarize yourself with reading/writing images from/to a file.
- Learn about some simple image processing algorithms.
- Improve your skills with manipulating arrays.
- Improve your skills with implementing constructors, destructors, copy-constructors, and overloading various operators.
- Learn to document and describe your programs.

Specifically, you have to write a package to implement the image data type using arrays. The image data type should allow you to:

(1) Read an image from a file.
(2) Save an image to a file.
(3) Get the info of an image.
(4) Set the value of a pixel.
(5) Get the value of a pixel.
(6) Extract a subimage from an image.
(7) Compute the average gray-level value of an image.
(8) Enlarge an image by some factor s.
(9) Shrink an image by some factor s.
(10) Reflect an image in the horizontal or vertical directions.
(11) Translate an image by some amount t.
(12) Rotate an image by some angle *theta*.
(13) Compute the sum of two images.
(14) Compute the difference of two images.
(15) Compute the negative of an image.

You will have to write a driver program which interacts with the above image data type and the user. The user should have the option to choose anyone of the above options *in any order*. Your program should be able to handle various user input errors.

Image type specification

```
class Image {
 public:
   constructor // default - no parameters
   constructor // with parameters
   destructor
   copy_constructor
   operator= //overload assignment
   getImageInfo
   getPixelVal
   setPixelVal
   inBounds
   getSubImage
   meanGray
   enlargeImage
   shrinkImage
   reflectImage
   translateImage
   rotateImage
   operator+ //overload addition for images
   operator- //overload subtraction for images
   reflect
   translate
   rotate
   negateImage
 private:
   int N; // no of rows
   int M; // no columns
   int Q; // no gray-level values
   int **pixelVal;
};
```

readImage(fileName, image) Reads in an image from a file. The pixel values are stored in an array and the image width, height, and number of gray-levels are recorded in the appropriate fields. A NOT-PGM exception should be raised if the image file to be read is not in PGM format (*image* is an object of type *ImageType*).

writeImage(fileName, image): Writes out an image to a file in the appropriate format (*image* is an object of type *ImageType*).

getImageInfo(noRows, noCols, maxVal): It returns the height (no of rows) of the image, the width (no of columns) of the image, and the max pixel value. (should be returned using "call by reference").

int getPixelVal(r, c): Returns the pixel value at *(r, c)* location. An OUT-OF-BOUNDS exception should be raised if *(r, c)* falls outside the image.

setPixelVal(r, c, value): Sets the pixel value at location *(r, c)* to *value*. An *OUT-OF-BOUNDS* exception should be raised if *(r, c)* falls outside the image.

bool inBounds(r, c): Returns true if the pixel *(r, c)* is inside the image.

getSubImage(ULr, ULc, LRr, LRc, oldImage): It crops a rectangular area within *oldImage*. Often, for image analysis, we want to investigate more closely a specific area within the image, called a Region of Interest (ROI). They are used to limit the extent of image processing operations to some small part of the image. The ROI is a rectangular area within the image, defined either by the coordinates of its upper-left (UL) and lower-right (LR) corners or by the coordinates of its upper-left corner and its dimensions. You can obtain the pixel coordinates of the UL and LR corners using *xv* by moving the cursor on the desired positions and by pressing the middle button (*Warning:* the first number displayed corresponds to *c* and the second to *r*). An *OUT-OF-BOUNDS* exception should be raised if UL or LR fall outside the image.

int meanGray(): Computes the average gray level value of an image (returns the results as an integer by truncating it).

enlargeImage(s, oldImage): Enlarges the input image *oldImage* by some *integer* factor *s*. Enlarging an image is useful for magnifying small details in an image. There are various ways to enlarge a given image. Here, we will use a simple method: to enlarge an image by a given factor *s*, we must replicate pixels such that each pixel in the input image becomes an *s*x*s* block of identical pixels in the output image. This technique is most easily implemented by iterating over pixels of the output image and computing the coordinates of the corresponding input image pixel.

shrinkImage(s, oldImage): Shrinks the input image *oldImage* by some *integer* factor *s*. Shrinking an image is useful, for example, to reduce a large image in size so that it fits on the screen. There are various ways to shrink a given image. Here, we will use a simple method: to shrink an image by a scale factor *s*, we must sample every *s*th pixel in the horizontal and vertical dimensions and ignore the others. Again, this technique is most easily implemented by iterating over pixels of the output image and computing the coordinates of the corresponding input image pixel.

reflectImage(flag, oldImage): Reflects the input image *oldImage* along the horizontal or vertical directions (determined by the boolean "flag"). Reflection along either direction can be performed by simply reversing the order of pixels in the rows or columns of the image.

translateImage(t, oldImage): Translates the input image *oldImage* by some amount *t*. The translation process can be performed with the following equations:

$$r' = r + t \tag{1}$$

$$c' = c + t \tag{2}$$

where *t* is an integer (note: the translation amounts in the horizontal and vertical directions can be different in general). There are some practical difficulties implementing translation using the above equations (see question 3 below).

rotateImage(theta, oldImage): Rotates the input image *oldImage* by some angle *theta*. The rotation process requires the use of the following equations:

$$r' = r\ cos(theta) - c\ sin(theta) \tag{3}$$

$$c' = r\ sin(theta) + c\ cos(theta) \tag{4}$$

where *theta* is the angle of rotation (positive values correspond to counterclockwise rotation). Although the above formula is the basis of rotation, it only gets you halfway there because it will rotate an image about point (0,0). In most cases, what we really want is to rotate an image about its center. The following equations rotate an image about its center:

$$r' = r_0 + (r - r_0)\ cos(theta) - (c - c_0)\ sin(theta) \tag{5}$$

$$c' = c_0 + (r - r_0)\ sin(theta) + (c - c_0)\ cos(theta) \tag{6}$$

There are some practical difficulties implementing rotation using (3)-(4) or (5)-(6). Let us consider what happens to pixel (0,100) after a 90 degrees rotation using equations (3)-(4):

$$r' = r\ cos(90) - c\ sin(90) = -100\ sin(90) = -100$$
$$c' = r\ sin(90) + c\ cos(90) = 0\ sin(90) = 0$$

In this case, the pixel moves to coordinates (-100,0). This is clearly a problem since pixels cannot have negative coordinates. Let's now consider what happens to pixel (50,0) after a 35 degrees rotation:

$$r' = r\ cos(35) - c\ sin(35) = 50\ cos(35) = 40.96$$
$$c' = r\ sin(35) + c\ cos(35) = 50\ sin(35) = 28.68$$

The coordinates calculated by the transformation equations are not integers, and therefore do not index a pixel in the output image.

The first problem can be resolved by testing coordinates to check that they lie within the bounds of the output image before attempting to copy pixels. A simple solution to the second problem is to find the nearest integers to r' and c' and use these as the coordinates of the transformed pixel.

Please note that the C++ math functions *cos()* and *sin()* require that the angle is given in *radians* (enter "man cos" or "man sin" from the command line to get a description of these functions). To convert degrees to radians, use the following formula:

$$\theta_{rad} = \frac{\theta_{deg}\ x\ \pi}{180.0} \tag{7}$$

To use these functions successfully, you need to include the following header file in your program:

#include <math.h>

Also, you need to "link" your program to the Math library (this can be done during compilation by appending *-lm* at the end of the command you are using to compile your program). There are some practical difficulties implementing rotation using the above equations (see question 4 below).

operator+: Computes the sum of two images. Addition is used to combine the information in two images. For example, you can implement simple "image morphing" using image addition (e.g., try adding together the following images from the image gallery: *person1.pgm*, *person2.pgm*, *person3.pgm*). If we add two 8-bit images, then pixels in the resulting image can have values in the range 0-510. One way to deal with this problem is by choosing a 16-bit representation (i.e., set Q=510) for the output image. Another way is to use the formula shown below:

$$O(r, c) = aI_1(r, c) + (1 - a)I_2(r, c) \tag{8}$$

where *a* is a constant in the interval [0,1] (addition is a special case when $a = 0.5$)

operator-: Computes the difference of two images. The main application of image subtraction is in *change detection*. If we make two observation of a scene and compute their difference, then changes will be indicated by pixels in the difference image which have non-zero values. If we subtract two 8-bit images, then pixels in the resulting image can have values between -255 and +255. This necessitates the use of 16 bit signed integers in the output image. However, the sign is usually unimportant and we just consider the absolute difference in which case we just need 8 bit integers:

$$O(r, c) = |I_1(r, c) - I_2(r, c)| \tag{9}$$

negateImage: Computes the negative of an image. This can be done by the following transformation:

$$O(r, c) = -I(r, c) + 255 \tag{10}$$

where *I(r,c)* is the input image (i.e., a grayscale image with 255 possible gray levels) and $O(r, c)$ is the output image.

Instructions

You should have three source code files: one for the application program containing the main function (*driver.cpp*), one header file *image.h* that specifies the image data type, and a file *image.cpp* that actually implements the image data type.

Describe the implementation of each function in detail. Each function should be discussed into a

separate section with the title of each section being the same as the name of the function. The sections should be clearly separated from each other.

Questions

Answering the follwoing questions does not require that you have prior knowledge of image processing. Just spend some time thinking about them and give us your best possible answers along with some justification. You do not have to do any extra coding regarding these questions but you are encouraged to do so. Interesting ideas which are implemented and demonstrated will get *extra credit* !! **Document in your report any extensions you might have made and let the TA know during the demo.**

1. (2pts extra) The algorithm you have to implement in this assignment to enlarge a given image is actually very simple. Can you think of a better approach? Compare your approach with the approach suggested in this assignment (i.e., list possible advantages/disadvantages).

2. (2pts extra) The algorithm you have to implement in this assignment to shrink a given image is actually very simple. Can you think of a better approach? Compare your approach with the approach suggested in this assignment (i.e., list possible advantages/disadvantages).

3. (2pts extra) There is a practical difficulty with the direct application of the translation equations (1) and (2). When translating, what is done with the "leftover" space? For example, if we move everything one row down, what do we put in the top row? One solution is to fill the top row with a constant value value (typically black (0) or white (255)). Can you think of a better solution?

4. (4pts extra) We have already discussed two practical difficulties associated with rotation: (a) the case where the transformed pixel coordinates fall outside the image and (b) the case where the transformed pixel coordinates are not integers. In the first case, we suggested simply ignoring the transformed pixel coordinates which fall outside the range. This is probably enough in most cases. In the second case, we suggested finding the nearest integer neighbors to r' and c'. This approach, however, will not produce a value for every pixel in the output image. In other words, it will produce numerous "holes" in the rotated image where no value was computed. Can you suggest a way for dealing with this problem?

5. (2pts extra) This question is regarding image subtraction. By examining the difference image, you will notice some rather small non-zero values in image areas where there is not really any change (e.g., subtract "backg1.pgm" from "backg2.pgm" and look at the values of the difference image). These differences are due to sensor noise, slight changes in illumination and various other factors. Can you suggest of a way to deal with this problem?