Hole Filling

Abstract

The goal of this task is to build a small image processing library that fills holes in images, along with a small command line utility that uses that library, and answer a few questions. The coding tasks must be implemented using one of the following languages: C++, Java, Kotlin, Objective-C or Swift. If you're not feeling proficient in any of these languages, please let us know in advance. Make sure to pay close attention to the design of your code, and not just to the completeness or efficiency of the code.

The library must support filling holes in grayscale images, where each pixel value is a `float` in the range `[0, 1]`, and hole (missing) values which are marked with the value `-1`.

The library should provide the ability to fill a hole in an image, based on the following algorithm:

Hole Filling Algorithm

Definitions

- I: the input image.
- $I(v)$: color of the pixel at coordinate $v \in \mathbb{Z}^2$.
- B: set of all the boundary pixel coordinates. A boundary pixel is defined as a pixel that is *connected* to a hole pixel, but is not in the hole itself. Pixels can be either 4- or 8-connected to the hole based on input. See this for more info.
- H: set of all the hole (missing) pixel coordinates. You can assume the hole pixels are 8-connected with each other.
- $w(v, u)$: weighting function which assigns a non-negative float weight to a pair of two pixel coordinates in the image.

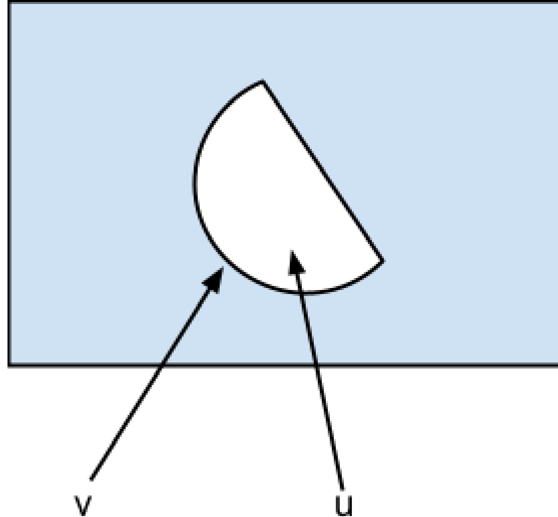

Figure 1: an image with a hole

Algorithm

Find the boundary $B$and the hole Hin a given image. For each $u \in H$, set its value to:

$$I(u) = \frac{\sum_{v \in B} w(u,v) \cdot I(v)}{\sum_{v \in B} w(u,v)}$$

Requirements

Default weighting function

The default weighting function for the algorithm is:

$$w_{z,\epsilon}(u,v) = \frac{1}{\|u - v\|^z + \epsilon}$$

where ϵ is a small float value used to avoid division by zero, and $||u - v||$ denotes the euclidean distance between u and v. The values $z,$ ϵ should be configurable.

Library

Write a library that provides implementation of the algorithm stated above. The library needs to support any arbitrary weighting function provided by the user and support both the 4- and the 8-boundary-connectivity options.

Command line utility

Write a command line utility that accepts an input image file(s), z, ϵ and connectivity type, fills the hole and writes the result to an image file.

The library should work with grayscale images, as noted above, with -1 marking a missing color. But the command line utility may accept (and write to) other image formats, and perform the required conversions. This might be more convenient (and easier) to implement, because the most popular image formats do not support `-1` values. You are free to choose what format (or formats) the command line utility accepts, and how to describe a hole in an image. Our suggestion is to accept two input RGB images, one for the image and one for a mask that defines the hole, and then merge the two images by converting RGB to grayscale and applying the mask to carve out the hole.

Note that in contrast to the library, the command line utility doesn't need to support an arbitrary weight function, only the default one with configurable z and ϵ.

General comments

1. For loading and saving the image files, and for conversion between different image formats, you may use OpenCV or any other similar library. You're allowed to use OpenCV's data structures to hold the image if you wish, but you're not allowed to use any algorithmic functionality (related to image processing) provided by it. Please note the Image I/O snippets section below for pointers on how to load an image file.
2. There is no need to create (build) the library separately from the command line utility, but the "library" code needs to be easily extractable to an actual library should you want to build one.
3. For simplicity, you can assume that every image has only a single hole.
4. There's no need to handle the corner case of holes that touch the boundaries of the image.

Questions

Please answer the following questions and add them to a separate file in your submission.

1. If there are m boundary pixels and n pixels inside the hole, what's the complexity of the algorithm that fills the hole, assuming that the hole and boundary were already found? Try to also express the complexity only in terms of n.
2. Describe an algorithm that approximates the result in $O(n)$ to a high degree of accuracy. Bonus: implement the suggested algorithm in your library in addition to the algorithm described above.
3. Bonus (hard!): Describe and implement an algorithm that finds the *exact* solution in $O(nlogn)$. In this section, feel free to use any algorithmic functionality provided by

external libraries as needed.

What you need to submit

1. Source code for the library and command line utility as described in the *Requirements* section. No need to submit binary libraries or compiled code.
2. Answers to the questions described in the *Questions* section
3. The result of applying the exercise on the attached image (Lenna.png) and mask (Mask.png) with 8-connectivity and the default weighting function using $z = 3$ and $\epsilon = 0.01$. The mask represents hole pixels as any pixel with an intensity value of less than `0.5`.

Good luck!

Appendix: Image I/O snippets

These are snippets to assist you with the part of the exercise that is less focused on. You can use them or alter them to your needs.

Java with BufferedImage

```
static void loadImage(String filePath) throws IOException {
  BufferedImage image = ImageIO.read(new File(filePath));
    for (int x = 0; x < image.getWidth(); x++) {
      for (int y = 0; y < image.getHeight(); y++) {
        Color rgb = new Color(image.getRGB(x, y));
        int grayscale = (int)(rgb.getRed() * 0.299 +
            rgb.getGreen() * 0.587 + rgb.getBlue() * 0.114);
        // Do something with grayscale...
    }
  }
}
```

Java with OpenCV

OpenCV is a native library, meaning you'll need to reference both a jar and a native library (DLL or SO) so you'll need to install OpenCV on your computer first - follow these steps.

```
import org.opencv.core.Core;
import org.opencv.core.Mat;
import org.opencv.imgcodecs.Imgcodecs;
import org.opencv.imgproc.Imgproc;

static void loadImage(String filePath) {
   System.loadLibrary(NATIVE_LIBRARY_NAME);
   Mat rgb = new Imgcodecs().imread(filePath);
   Mat image = new Mat();
   Imgproc.cvtColor(rgb, image, Imgproc.COLOR_RGB2GRAY);
   for (int x = 0; x < image.width(); x++) {
       for (int y = 0; y < image.height(); y++) {
           double grayscale = image.get(y, x)[0];
           // Do something with grayscale...
       }
   }
}
```

C++ with OpenCV

you'll need to install OpenCV on your computer first
https://learnopencv.com/install-opencv-on-windows/

```
#include <opencv2/opencv.hpp>
#include <string>

static void loadImage(std::string path) {
  image = cv::imread(path, IMREAD_GRAYSCALE);
  image.convertTo(image, CV_32FC1, 1.0 / 255.0);
  for (int i = 0; i < image.rows; ++i) {
    for (int j = 0; j < image.cols; ++j) {
      float grayscale = image.at<float>(i, j);
      // Do something with grayscale...
    }
  }
}
```

Swift with CGContext

```swift
import CoreGraphics

func loadImage(from path: String) throws {
  guard let image = NSImage(contentsOfFile: path) else {
    // ...
  }

  var imageRect = CGRect(x: 0, y: 0, width: image.size.width,
                         height: image.size.height)
  guard let cgImage = image.cgImage(forProposedRect: &imageRect,
                                    context: nil, hints: nil) else {
    // ...
  }

  var pixelData = [UInt8](repeating: 0, count: cgImage.width * cgImage.height)
  let colorSpace = CGColorSpaceCreateDeviceGray()
  guard let context = CGContext(data: &pixelData, width: cgImage.width,
                                height: cgImage.height, bitsPerComponent: 8,
                                bytesPerRow: cgImage.width, space: colorSpace,
                                bitmapInfo: 0) else {
    // ...
  }
  context.draw(cgImage, in: CGRect(x: 0, y: 0, width: cgImage.width,
                                   height: cgImage.height))

  for grayscale in pixelData {
    // Do something with grayscale
  }
}
```