

COMP 204

Introduction to image analysis with scikit-image (part one)

Yue Li

based on slides from Mathieu Blanchette, Christopher J.F.
Cameron and Carlos G. Oliver

Image processing and analysis in Python

Goal: Process and analyze digital images.

- ▶ Very useful for processing microscopy images, medical imaging, etc.
- ▶ Closely linked with machine learning for image analysis

scikit-image module or (skimage)

- ▶ image processing module in Python
- ▶ holds a wide library of image processing algorithms: filters, transforms, point detection
- ▶ API
 - ▶ <http://scikit-image.org/docs/dev/api/api.html>

Our mascot for today




What's an image in Python?

An image is stored as a NumPy ndarray (n-dimensional array).

- ▶ ndarrays are easier and more efficient than using 2-dimensional lists as we've seen before.

A color image with R rows and C columns is

- ▶ represented as a 3-dimensional ndarray of dimensions $R \times C \times 3$
- ▶ element at position (i, j) of the array corresponds to the RGB value at row i and column j
- ▶ each pixel is represented by 3 numbers, each between 0 and 255: Red, Green, Blue



Reading an image into memory

We'll start with an example image using the **io module**

- ▶ basic I/O submodule of scikit-image
- ▶ API:

<http://scikit-image.org/docs/dev/api/skimage.io.html>

```
1 import skimage.io as io
2 import matplotlib.pyplot as plt
3
4 # read image into memory
5 image = io.imread("monkey.jpg")
6
7 plt.imshow(image)
8 plt.show()
9
10 # print top-left pixel RGB values
11 print(image[0,0]) # a white pixel in the monkey image
12 # prints: [255 255 255]
13 # write image to disk
14 io.imsave("monkey_copy.jpg", image)
```

Playing with an image

```
5 image = io.imread("monkey.jpg")
6
7 n_row, n_col, n_colours = image.shape
8 print(n_row, n_col, n_colours) # prints (1362, 2048, 3)
9
10 # print pixel value at row 156 and column 293
11 pixel = image[156,292]
12 print(pixel) # prints [112 158 147] = green-bluish color
13
14 # Create a pink rectangle between rows 700-800
15 # and column 1000-1200
16 for i in range(700,800):
17     for j in range(1000,1200):
18         image[i,j] = (255,0,255)
19
20 # this is equivalent to the following code
21 # i.e., set every element in the array to (255,0,255)
22 #image[700:800,1000:1200] = (255,0,255)
23
24 plt.imshow(image)
25 plt.show()
26 io.imsave("monkey_bar.jpg", image)
```

A face-masked image



Let's try a more "refined" face mask by modifying the code

Code for a more refined mask

```
1 import skimage.io as io
2 import matplotlib.pyplot as plt
3
4 # read image into memory
5 image = io.imread("monkey.jpg")
6
7 # find out where the monkey face is in the image
8 plt.imshow(image)
9 plt.show()
10
11 # a more "refined" face mask (713,785),(1070,1150)
12 image[713:785,1070:1150] = (255,0,255)
13
14 plt.imshow(image)
15 plt.show()
16 io.imsave("monkey_refined_mask.jpg", image)
```

A more refined face-masked image



Creating the negative of an image

```
1 import skimage.io as io
2 import skimage.color as color
3 import matplotlib.pyplot as plt
4
5 # read image into memory
6 image = io.imread("monkey.jpg")
7
8 # Create the negative of an image
9 for i in range(image.shape[0]):
10     for j in range(image.shape[1]):
11         for c in range(3):
12             image[i,j,c] = 255-image[i,j,c]
13
14 # we could just have written:
15 #image = 255 - image
16
17 plt.imshow(image)
18 plt.show()
19 io.imsave("monkey_negative.jpg", image)
```

The “negative” image




Flipping the image horizontally

There is a bug in this code. Can you find and fix it?

```
1 import skimage.io as io
2 import skimage.color as color
3 import matplotlib.pyplot as plt
4
5 # read image into memory
6 image = io.imread("monkey.jpg")
7 n_row, n_col, colours = image.shape
8
9 # Flip the image horizontally
10 for i in range(0,n_row):
11     for j in range(0,int(n_col/2)):
12         image[i,j] = image[i, n_col-j-1]
13
14 plt.imshow(image)
15 plt.show()
16 io.imsave("monkey_flipped_wrong.jpg",image)
```

Incorrectly flipped image (a gemini monkey? just kidding)



Flipping the image horizontally (3 correct ways)

```
1 import skimage.io as io
2 import skimage.color as color
3 import matplotlib.pyplot as plt
4 import numpy as np
5
6 # read image into memory
7 image = io.imread("monkey.jpg")
8 n_row, n_col, colours = image.shape
9
10 # Flip the image horizontally
11 for i in range(0,n_row):
12     for j in range(0,int(n_col/2)):
13         t = image[i,j].copy()
14         image[i,j] = image[i, n_col-j-1]
15         image[i, n_col-j-1] = t
16
17 # this is equivalent to:
18 #image = image[:,::-1]
19
20 # this is also equivalent to:
21 #image = np.flip(image, 1)
22
23 plt.imshow(image)
24 plt.show()
25 io.imsave("monkey_flipped_right.jpg",image)
```

Correctly flipped image



How to flip an image up side down?



Combining images

Since images are just numerical arrays, we can easily combine them. Example:

Create an image that is the average of monkey and tiger.



Combining images

```
1 import skimage.io as io
2 import matplotlib.pyplot as plt
3 import numpy as np
4 from skimage.transform import resize
5
6 monkey = io.imread("monkey.jpg")
7 tiger = io.imread("tiger.jpg")
8
9 #resize images to 500x1000 pixels
10 monkey_resized = resize(monkey, (500, 1000))
11 tiger_resized = resize(tiger, (500, 1000))
12
13 combined = np.zeros((500,1000,3))
14 for i in range(500):
15     for j in range(1000):
16         for c in range(3):
17             combined[i,j,c]=monkey_resized[i,j,c]/2 + \
18                 tiger_resized[i,j,c]/2
19
20 # this is equivalent to:
21 #combined = monkey_resized/2 + tiger_resized/2
22
23 plt.imshow(combined)
24 plt.show()
25 io.imsave("combined.jpg",combined)
```


Combining images (monkey/2+tiger/2)



crop and combe images

```
1 import skimage.io as io
2 import matplotlib.pyplot as plt
3 from skimage import img_as_float
4
5 monkey = io.imread("monkey.jpg")
6 tiger = io.imread("tiger.jpg")
7
8 print(monkey.shape) # (1362, 2048, 3)
9 print(tiger.shape) # (697, 1400, 3)
10
11 monkey_height, monkey_width, n_colours = monkey.shape
12 tiger_height, tiger_width, n_colours = tiger.shape
13
14 monkey_cropped =
15     ↪ monkey[monkey_height-tiger_height:monkey_height,
16             monkey_width-tiger_width:monkey_width]
17
18 print(monkey_cropped.shape) # (697, 1400, 3)
19 print(tiger.shape) # (697, 1400, 3)
20
21 combined = monkey_cropped/2 + tiger/2
22 combined = img_as_float(combined/255)
23
24 plt.imshow(combined)
25 plt.show()
26 io.imsave("cropped_and_combined.jpg", combined)
```

cropped and combined images




Switching patches between two images

```
8  #plt.imshow(monkey)
9  #plt.show()
10 print(monkey.shape)
11 monkey_head = monkey[618:774,1017:1176]
12 #plt.imshow(monkey_head)
13 #plt.show()
14
15 #plt.imshow(monkey)
16 #plt.show()
17 print(tiger.shape)
18 tiger_head = tiger[72:257,606:762]
19 #plt.imshow(tiger_head)
20 #plt.show()
21
22 print(monkey_head.shape)
23 print(tiger_head.shape)
24
25 tiger_head_resized = resize(tiger_head, (774-618,
   ↪ 1176-1017))
26 monkey_head_resized = resize(monkey_head, (257-72,
   ↪ 762-606))
27
28 monkey[618:774,1017:1176] = tiger_head_resized*255
29 tiger[72:257,606:762] = monkey_head_resized*255
```

A monkey with a tiger head




A tiger with a monkey head



Color separation colors

```
1 import skimage.io as io
2
3 image = io.imread("monkey.jpg")
4
5 red, green, blue = image.copy(), image.copy(),
  ↪ image.copy()
6 red[:, :, (1,2)] = 0 # NumPy indexing
7 green[:, :, (0,2)] = 0
8 blue[:, :, (0,1)] = 0
9 io.imsave("monkey_red.jpg", red)
10 io.imsave("monkey_green.jpg", green)
11 io.imsave("monkey_blue.jpg", blue)
```



red intensity



green intensity



blue intensity

Shifting colors

```
1 import skimage.io as io
2 import skimage.color as color
3 import matplotlib.pyplot as plt
4 import numpy as np
5
6 image = io.imread("monkey.jpg")
7 n_row, n_col, colors = image.shape
8
9 # create a blank image
10 new_image = np.zeros( (n_row, n_col, 3), dtype=np.uint8)
11
12 # assemble a new image made of shifted colors
13 # blue is shifted right by 100 pixels
14 # green is shifted up by 100 pixels
15 for i in range(n_row):
16     for j in range(n_col):
17         new_image[i,j,0] = image[i,j,0] # keep red
18         if i>=100:
19             new_image[i,j,1]=image[i-100,j,1] # move
20                 ↵ green
21         if j>=100:
22             new_image[i,j,2]=image[i,j-100,2] # move blue
23
24 plt.imshow(new_image)
25 plt.show()
26 io.imsave("monkey_shifted.jpg",new_image)
```

Color shifted image



Grayscaleing

Many image processing algorithms assume a 2D matrix

- ▶ not an image with a third dimension of color

To bring the image into two dimensions

- ▶ we need to summarize the three colors into a single value
- ▶ this process is more commonly know as **grayscaleing**
- ▶ where the resulting image only holds intensities of gray
 - ▶ with values between 0 and 1

skimage submodule **color** has useful functions for this task

- ▶ API:

<http://scikit-image.org/docs/dev/api/skimage.color.html>

Grayscaleing images

```
1 from skimage.color import rgb2gray
2 import skimage.io as io
3 import matplotlib.pyplot as plt
4
5 # read image into memory
6 image = io.imread("monkey.jpg")
7 # convert to grayscale
8 gray_image = rgb2gray(image)
9 plt.imshow(gray_image, cmap='gray')
10 plt.show()
11 io.imsave("monkey_grayscale.jpg", gray_image)
12
13 print(image[0,0]) # a white pixel in RGB
14 # prints: [255 255 255]
15 print(gray_image[0,0]) # a white pixel in grayscale
16 # prints: 1.0
```

Grayscale image



Binarizing images

```
1 import skimage.io as io
2 import matplotlib.pyplot as plt
3 from skimage.color import rgb2gray
4 import numpy as np
5
6 image = io.imread("monkey.jpg")
7 gray_image = rgb2gray(image)
8 #print(gray_image[0,0])
9 # prints: 1.0
10 binary_image = np.where(gray_image >
    ↪ np.mean(gray_image),1.0,0.0)
11
12 plt.imshow(binary_image, cmap='gray')
13 plt.show()
14 io.imsave("monkey_binary.jpg",binary_image)
15 print(binary_image[0,0]) # a white pixel in the
    ↪ binary image
16 # prints: 1.0
```

`numpy.where(condition[, x, y]):` Return elements chosen from `x` or `y` depending on the condition.

A binary image

