



Bitmapped Images

Multimedia Techniques & Applications

Yu-Ting Wu

Outline

- Overview
- Image compression
- Image manipulation
- Geometrical transformations

Overview

- Record the value of every pixel in the image
- **Image size** is the main cost for the simplicity
- Images created from external devices are usually in a bitmapped fashion
 - Digital cameras
 - Scanners



digital camera



scanner

Resolution

- A measure of **how finely** a device approximates continuous images using finite pixels
 - Closely related to sampling rates
- Two ways of specifying resolution
 - **Printers and scanners: number of dots per unit**
 - Dots per inch (dpi)
 - Ex: consumer printer (600 dpi), book production (1200 – 2700 dpi), scanners (300 dpi – 3600 dpi)
 - **Video: size of a frame measured in pixels**
 - Ex: 640 x 480, 768 x 576
 - Can translate into the form of dpi if you know the physical dimension of the display device

Image Compression


Image Compression

- **Motivation**

- Faithfully store all pixel values of an image takes lots of memory space
- Human eyes can tolerate some minor errors in images
 - Digital representation is an approximation itself


- **Assumption**

- Images are usually **smooth** and have some **spatial coherence**




Compression Methods

- Spend some computation efforts to earn saving in space
- The effectiveness depends on the content of the compressed image
 - **Image size can become bigger after applying compression**
 - Definitely true, otherwise, any data can be compressed into one byte



128 bytes → 6 bytes
for a row (RLE)




128 bytes → 256 bytes
for a row (RLE)

Compression Methods (cont.)


- **Lossless compression**

- No information will lose during a compression/decompression cycle
- Ex: run-length encoding (RLE), variable-length coding




- **Lossy compression**

- Discard some information during the compression process and the information can **never** be recovered
- Ex: JPEG



Run-Length Encoding (RLE)

- The simplest compression technique
- Each time store a value, followed by a **count** to indicate a number of consecutive pixels of that value
- Example
 - RLE for row1: 128 128
 - 128 bytes (raw) v.s. **2 bytes** (using short)
 - RLE for row2: 128 32 0 64 128 32
 - 128 bytes (raw) v.s. **6 bytes** (using short)










Huffman Coding

- The best-known variable length coding
- **Lossless** compression
- Example:

Assume an 8 x 8 image containing 6 different pixel intensities

We can count their occurrence:







						
intensity	20	60	100	140	180	220
occurrence	5	6	25	16	9	3




Huffman Coding (cont.)


• Algorithm

- Build a Huffman tree
 - Sort the occurrence of intensity

					
20	60	100	140	180	220
5	6	25	16	9	3









- Merge the two with the smallest occurrence, and sort again




Huffman Coding (cont.)

• Algorithm

- Build a Huffman tree
 - Keep doing ...







					
20	60	100	140	180	220
5	6	25	16	9	3




Huffman Coding (cont.)

• Algorithm

- Build a Huffman tree
 - Keep doing ...


					
20	60	100	140	180	220
5	6	25	16	9	3



Huffman Coding (cont.)

• Algorithm

- Build a Huffman tree
 - Keep doing ...








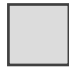
20	60	100	140	180	220
5	6	25	16	9	3

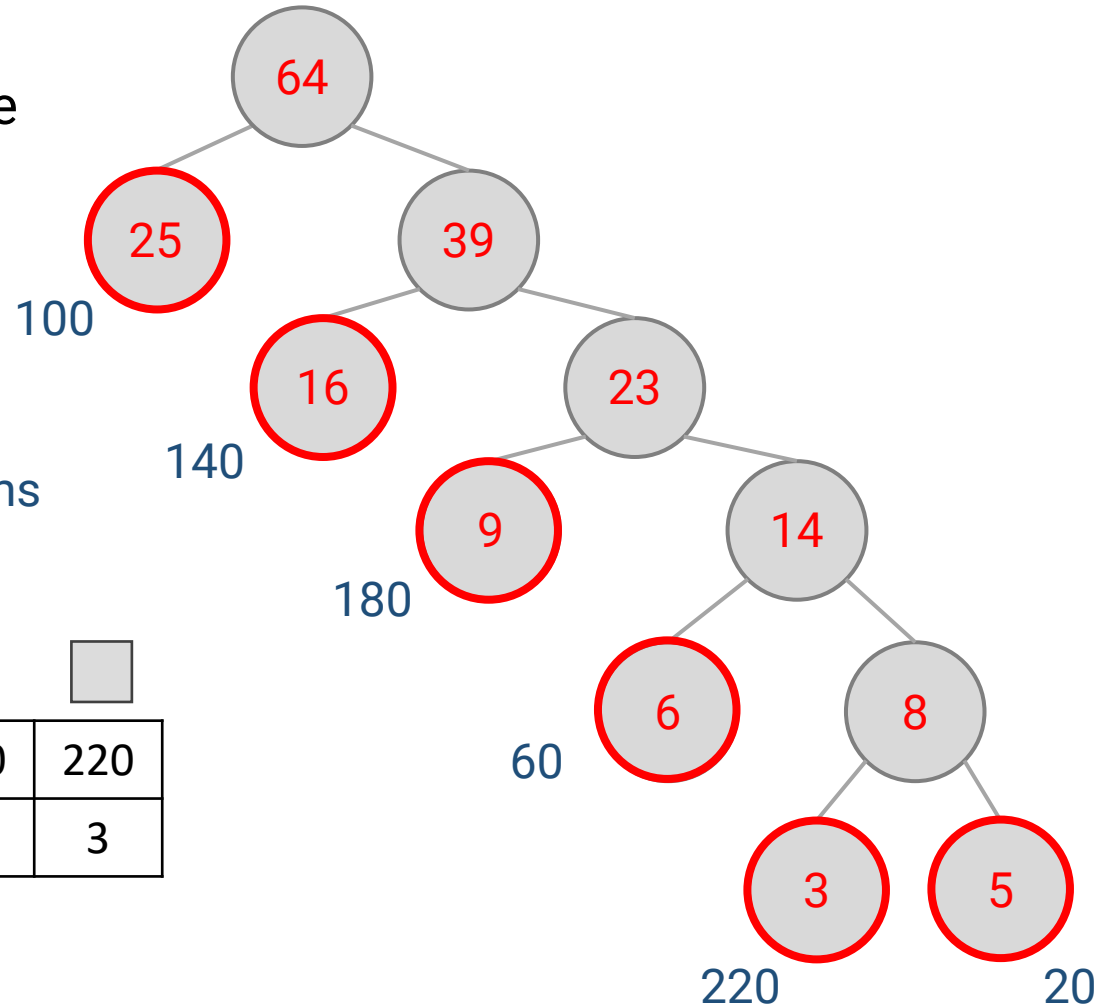
Huffman Coding (cont.)

• Algorithm

- Build a Huffman tree
 - Keep doing ...

Once we have done, the leaf nodes are the initial data items


					
20	60	100	140	180	220
5	6	25	16	9	3









Huffman Coding (cont.)

• Algorithm

- Build a Huffman tree
- Label the code (from root)



					
20	60	100	140	180	220
5	6	25	16	9	3

Huffman Coding (cont.)

• Algorithm

- Build a Huffman tree
- Label the code (from root)
- We will obtain

100: 0


140: 10

180: 110

60: 1110

220: 11110

20: 11111




Average number of bit per data:

$$1 \cdot (25/64) + 2 \cdot (16/64) + 3 \cdot (9/64) + 4 \cdot (6/64) + 5 \cdot (3/64) + 5 \cdot (5/64) \\ = \mathbf{2.31} \text{ (compared to raw, 8 bit, and naïve encoding, 3 bit)}$$

JPEG Compression

- JPEG is the most important **lossy** compression technique, which stands for ***Joint Photographic Experts Group***
 - Related file formats: *.jpg / *.jpeg / *.jpe / *.jfif / *.jfi / *.jif
- It works because image data can tolerate a certain amount of data loss

JPEG Compression (cont.)



• RGB → YCbCr

- People are more sensitive to intensity (Y) and less sensitive to color (Cb, Cr)
- Cb and Cr are lower frequency and have more spatial coherence
- Compress Cb and Cr; while keep Y as it is



RGB



Y




Cb




Cr

JPEG Compression (cont.)




- **Divide into 8x8 blocks (for Cb & Cr)**

- The entire image is too difficult to compress
- Small image block has higher coherence



JPEG Compression (cont.)




• Discrete Cosine Transform (DCT)



- A method for transforming a waveform into its frequency domain
- The DCT of an image block is the coefficients of different cosines of the image block




image




DFT result



DCT result




JPEG Compression (cont.)



- **Quantization**


- Human are less sensitive to high-frequency signal
- Use fewer bits for high-frequency signals in the DCT result and vice versa
- **This step is the reason of lossy compression**
- After this step, many components will end up with zero coefficients

JPEG Compression (cont.)




- **Zigzag ordering**

- For later Huffman encoding
- Result in a longer zero sequence



JPEG Compression (cont.)



- **RLE / Huffman encoding**

- Different strategy for DC and AC term
- The DC components of different blocks are encoded using Huffman algorithm
- The AC components within a block are encoded using RLE



JPEG Compression (cont.)

- The decompression of JPEG data is done by reversing the compression process
- We can control the degree of compression by altering the amount of quantization
- JPEG compression usually achieves very high compression rate for natural images (5% of the original size)

Image Manipulation

Image Manipulation

- **Motivations**
 - Correct deficiencies in an image (e.g., noise, red-eye)
 - Create images that are difficult or impossible to make naturally (e.g., glow)
- Type of image manipulations
 - Pixel **point** processing
 - Pixel **group** processing

Pixel Point Processing

- Compute a pixel's new value solely on the basis of its old value

mapping function

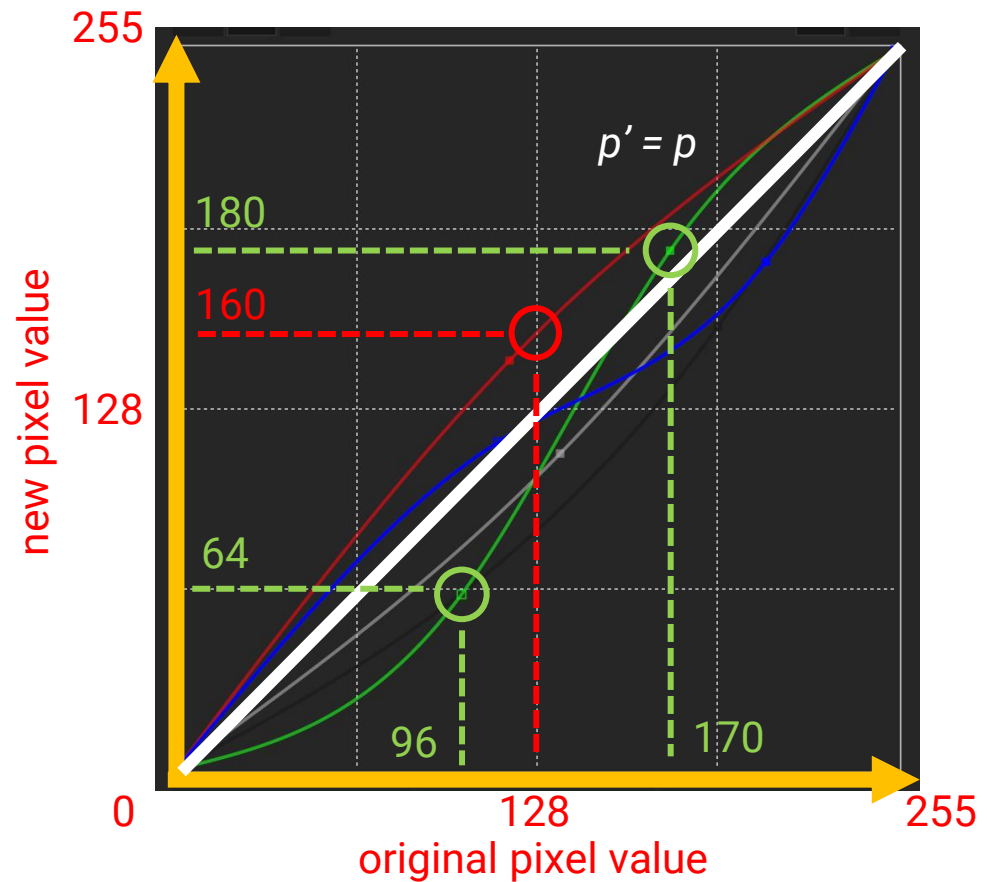
$$p' = f(p)$$

- Some examples
 - Adjustment of brightness
 - Adjustment of contrast
 - Change the black and white levels



Color Curve

- The operations can be considered generally as altering the mapping function f
- **Color curve**





Histogram

- An approximate representation of the **distribution** of numerical data




Intensity Image




Pixel Group Processing

- Compute each pixel's new value as a function not just of its old value, but also of **the values of neighboring pixels**
- Usually related to **filtering**
 - For a pixel of an image, specify a two-dimensional array of weights of its neighbors
 - Several types of filters
 - Smoothing
 - Sharpening
 - Detecting edge



Box Filter

- Each neighbor has the same weight



Box Filter (cont.)

convolution mask
convolution kernel

↓

$1/9$	$1/9$	$1/9$
$1/9$	$1/9$	$1/9$
$1/9$	$1/9$	$1/9$

filter

0	0	9	9
0	0	9	9
0	0	9	9
0	0	9	9

image (signal)

$$(1/9*0+1/9*0+1/9*0+1/9*0+1/9*0+1/9*0+1/9*0+1/9*9+1/9*9+1/9*9)$$

0	3	6	9
0	3	6	9
0	3	6	9
0	3	6	9


filtered image (signal)



Input $f(x,y)$

$$* \frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} =$$


Convolution operator,
not multiplication!




Output $g(x,y)$

Box Filter (cont.)

- 1D visualization of kernel weight



box filter



gaussian filter



Gaussian Filter

- The weight of neighbor falls exponentially with its distance to the filtered pixel
 - Standard deviation σ controls the speed of decreasing

$$g(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

distance
standard deviation

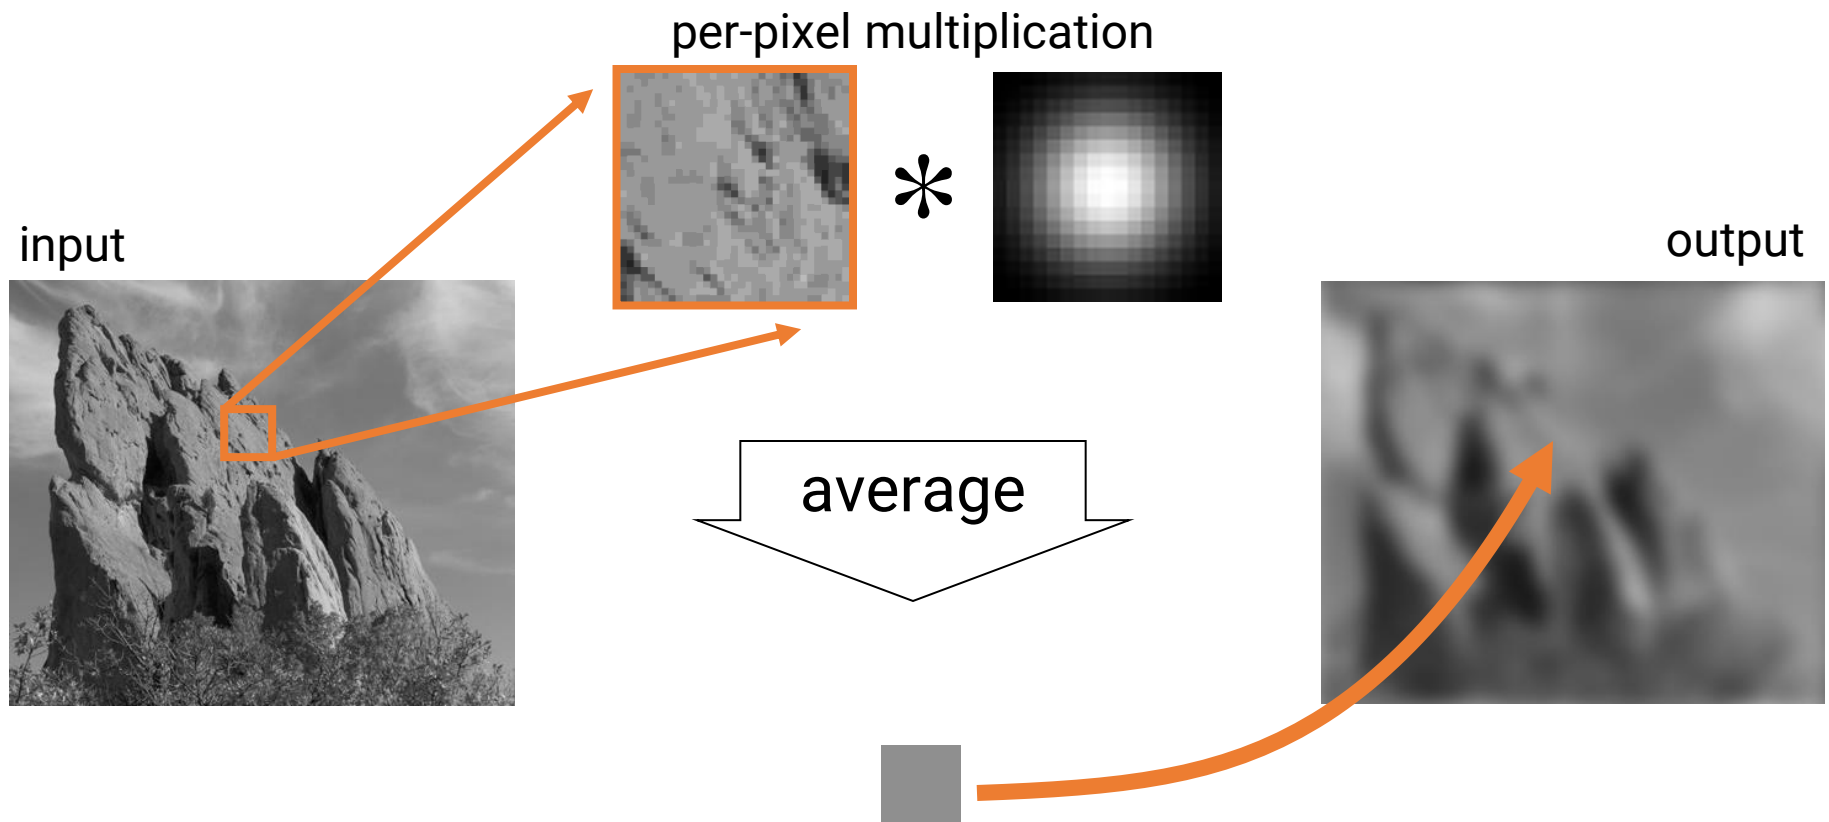
normalization
(sum of weight = 1)



$$\frac{1}{273}$$

1	4	7	4	1
4	16	26	16	4
7	26	41	26	7
4	16	26	16	4
1	4	7	4	1

Gaussian Filter (cont.)




Box Filter v.s. Gaussian Filter


original image



box filter



gaussian filter




Bilateral Filter

- Properties of Gaussian filter



$$g(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{\underbrace{2\sigma^2}_{\text{range of the filter}}}}$$



input

small σ 

limited smoothing

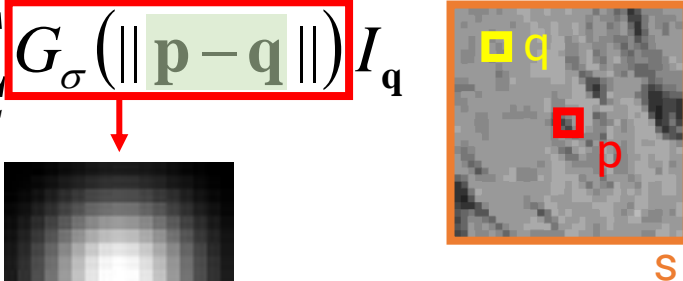
large σ 

strong smoothing

Bilateral Filter (cont.)

- Problems of Gaussian filter
 - Does smooth images
 - But smooths too much: **edges are blurred**
 - Only spatial distance matters
 - No edge term

spatial distance

$$GB[I]_p = \sum_{q \in S} G_{\sigma}(\|p - q\|) I_q$$


$$g(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

input




output



Bilateral Filter (cont.)

- Problems of Gaussian filter
 - Same Gaussian kernel everywhere



Bilateral Filter (cont.)

- Combine another Gaussian weight computed by **intensity difference (edge preserving)**

$$GB[I]_p = \sum_{q \in S} G_{\sigma}(\|\mathbf{p} - \mathbf{q}\|) I_q$$

spatial distance




$$BF[I]_p = \frac{1}{W_p} \sum_{q \in S} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(\|I_p - I_q\|) I_q$$

spatial distance range (intensity) distance
NEW!

Bilateral Filter (cont.)


- Visualization

$$BF [I]_p = \frac{1}{W_p} \sum_{q \in S} \underbrace{G_{\sigma_s} (\| \mathbf{p} - \mathbf{q} \|)}_{\text{Spatial Similarity}} \underbrace{G_{\sigma_r} (\| I_p - I_q \|)}_{\text{Range Similarity}} I_q$$



Bilateral Filter (cont.)

- Parameters

$$BF [I]_{\mathbf{p}} = \frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_s} (\| \mathbf{p} - \mathbf{q} \|) G_{\sigma_r} (| I_{\mathbf{p}} - I_{\mathbf{q}} |) I_{\mathbf{q}}$$


- Spatial sigma σ_s** : spatial extent of the kernel, size of the considered neighborhood
- Range sigma σ_r** : “minimum” amplitude of an edge

Bilateral Filter (cont.)

- Parameters

$$\sigma_s = 2$$



$$\sigma_r = 0.25$$





$$\sigma_r = \infty$$



 (Gaussian blur)



$$\sigma_s = 6$$




$$\sigma_s = 18$$



input

Bilateral Filter Application (cont.)

- Denoising




noisy input




Median Filter 5x5

Bilateral Filter Application (cont.)

- Denoising



noisy input



Bilateral Filter 7x7

Bilateral Filter Application (cont.)

- Denoising



noisy input



Gaussian Filter



Bilateral Filter

Sobel Filter


- Negative weights are commonly used for edge detection

-1	0	+1
-2	0	+2
-1	0	+1

Gx

+1	+2	+1
0	0	0
-1	-2	-1

Gy



Photographic Style Transfer

- Two-scale Tone Management for Photographic Look, Bae et al. SIGGRAPH 2006



Photographic Style Transfer (cont.)

- Motivation



Ansel Adams
Clearing Winter Storm



An Amateur Photographer

Photographic Style Transfer (cont.)

- Motivation



Ansel Adams
Clearing Winter Storm




style transferred result
(an imitation of Ansel Adams)

Photographic Style Transfer (cont.)

• Observation

- Different photographers have different tonal styles
 - Global contrast



Ansel Adams




Kenro Izu

Photographic Style Transfer (cont.)

• Observation

- Different photographers have different tonal styles
 - Global contrast
 - Local contrast



Ansel Adams



Kenro Izu

Photographic Style Transfer (cont.)

- **Goal**
 - Transfer look between photographers



input image



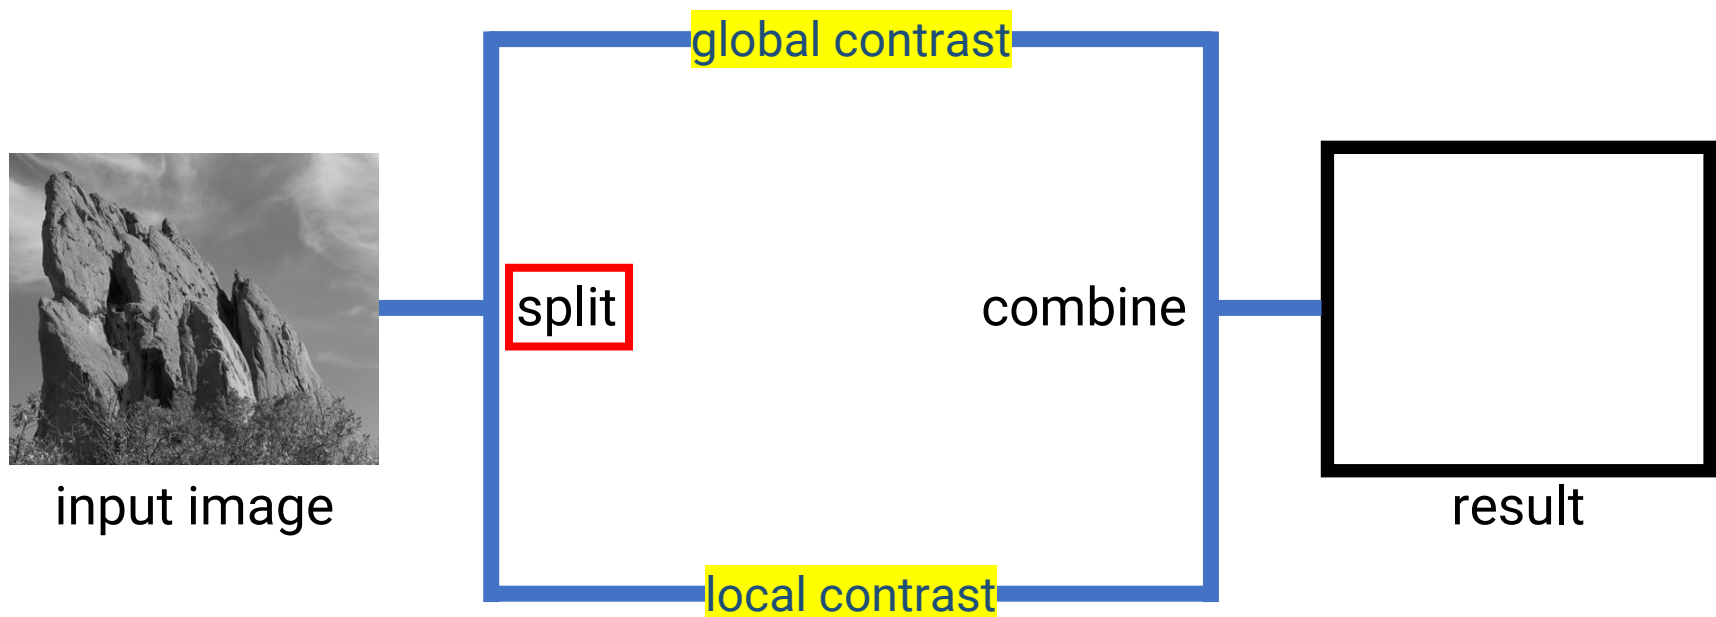
model



result

Photographic Style Transfer (cont.)


- Algorithm



Photographic Style Transfer (cont.)

- **Algorithm**

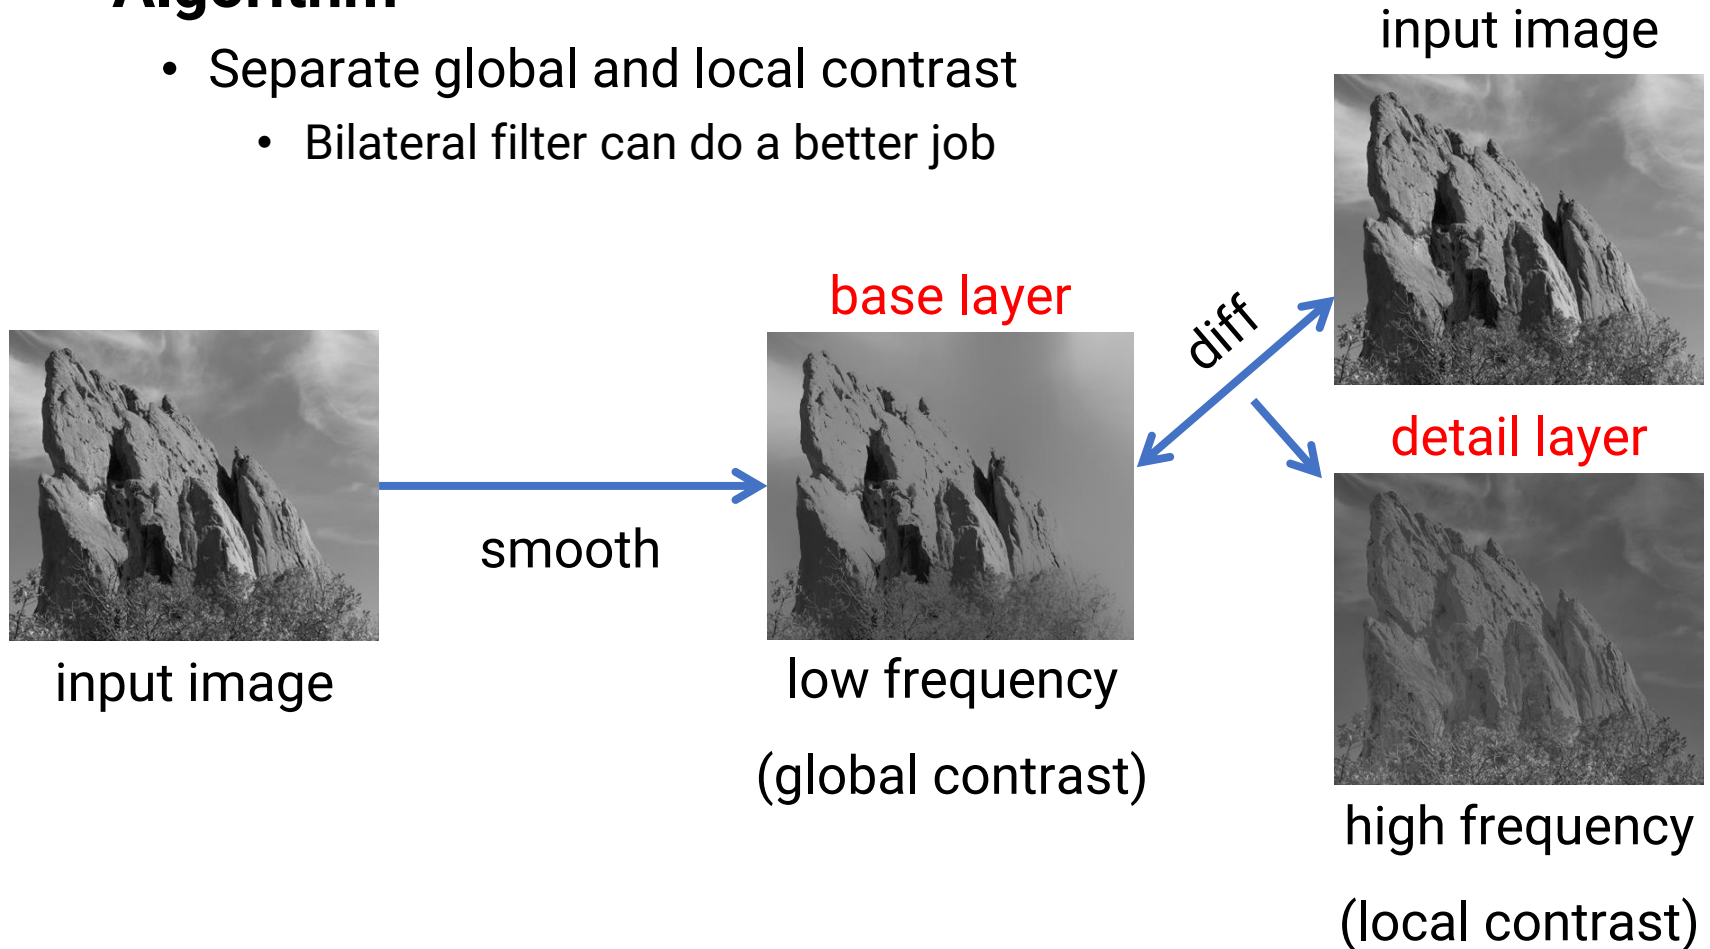
- Separate global and local contrast
 - Gaussian filter produces blurring and halos



Photographic Style Transfer (cont.)

- **Algorithm**

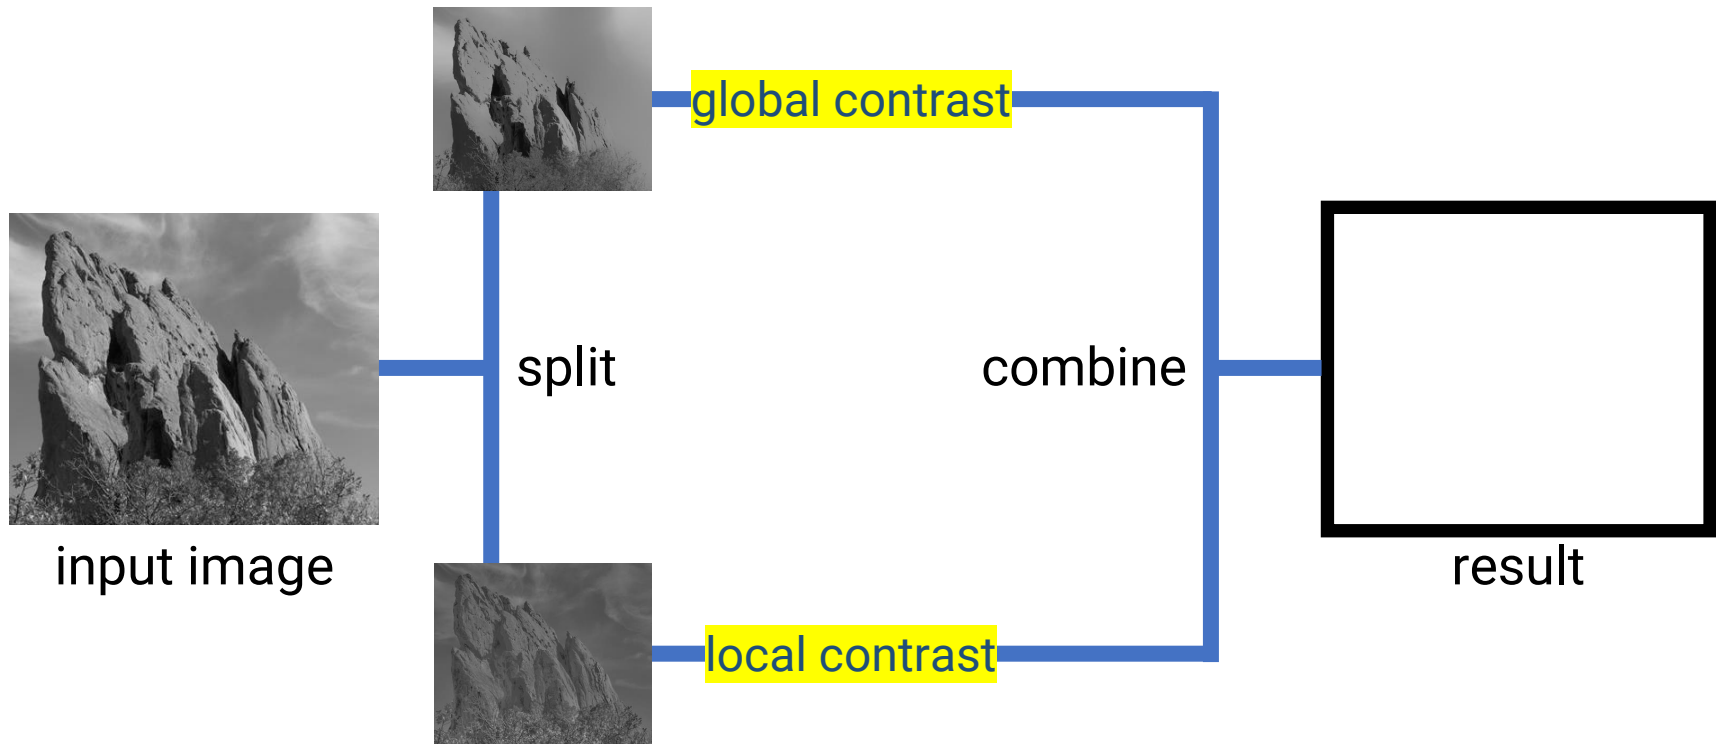
- Separate global and local contrast
 - Bilateral filter can do a better job



Photographic Style Transfer (cont.)

- **Algorithm**


- Separate global and local contrast



Photographic Style Transfer (cont.)

- **Algorithm**


- Adjust global contrast by histogram matching



Photographic Style Transfer (cont.)

- **Algorithm**


- Adjust local contrast by uniformly scaling



Photographic Style Transfer (cont.)

- **Algorithm**

- Sometimes the local contrast is not uniform



smooth region

⇒ low textureness

textured region


⇒ high textureness

Photographic Style Transfer (cont.)

- **Algorithm**
 - Textureness computation



input

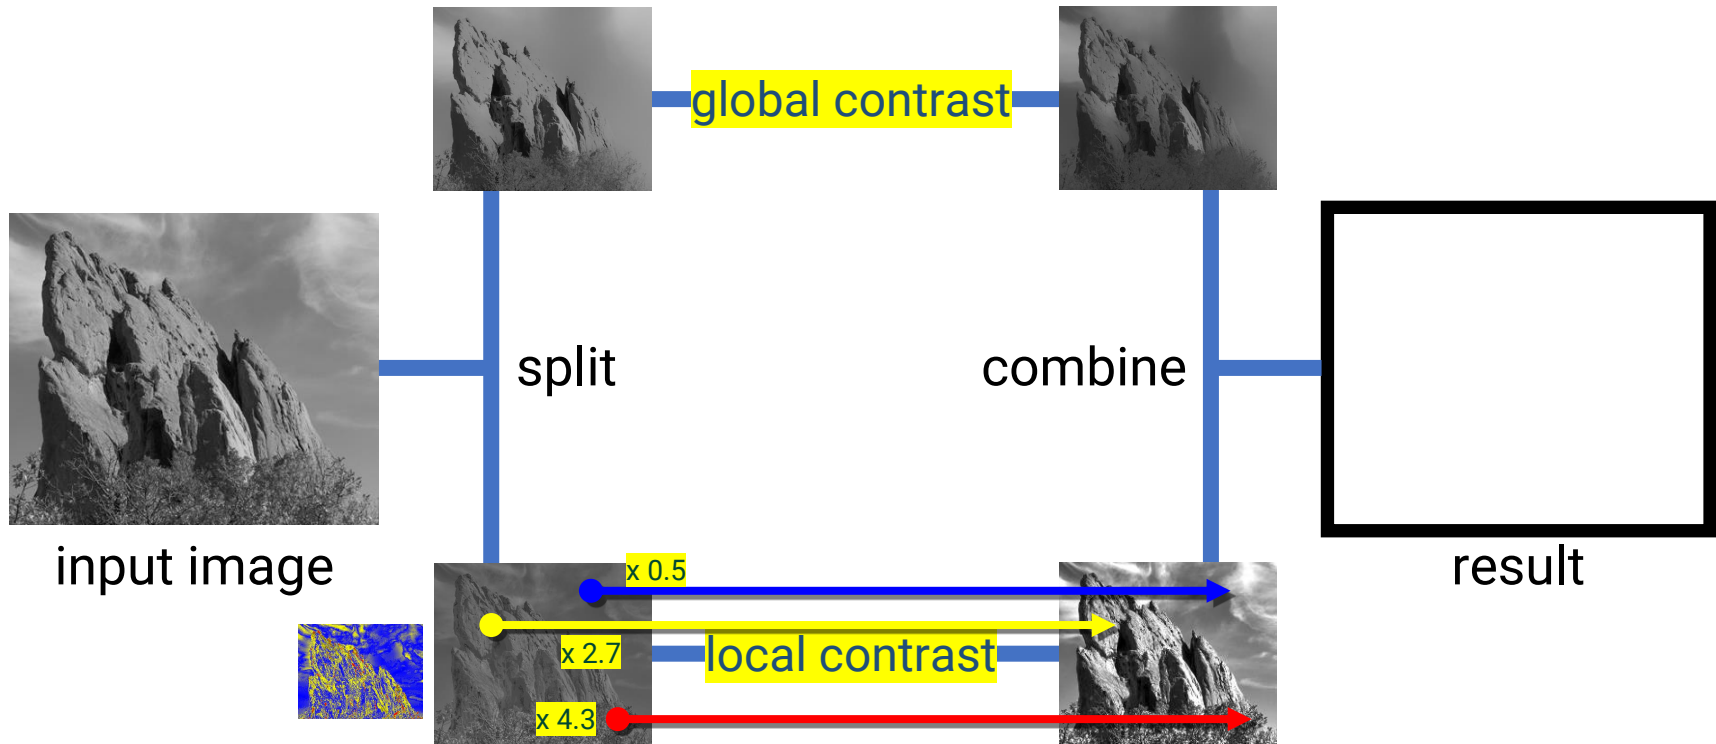


textureness

Photographic Style Transfer (cont.)

- **Algorithm**

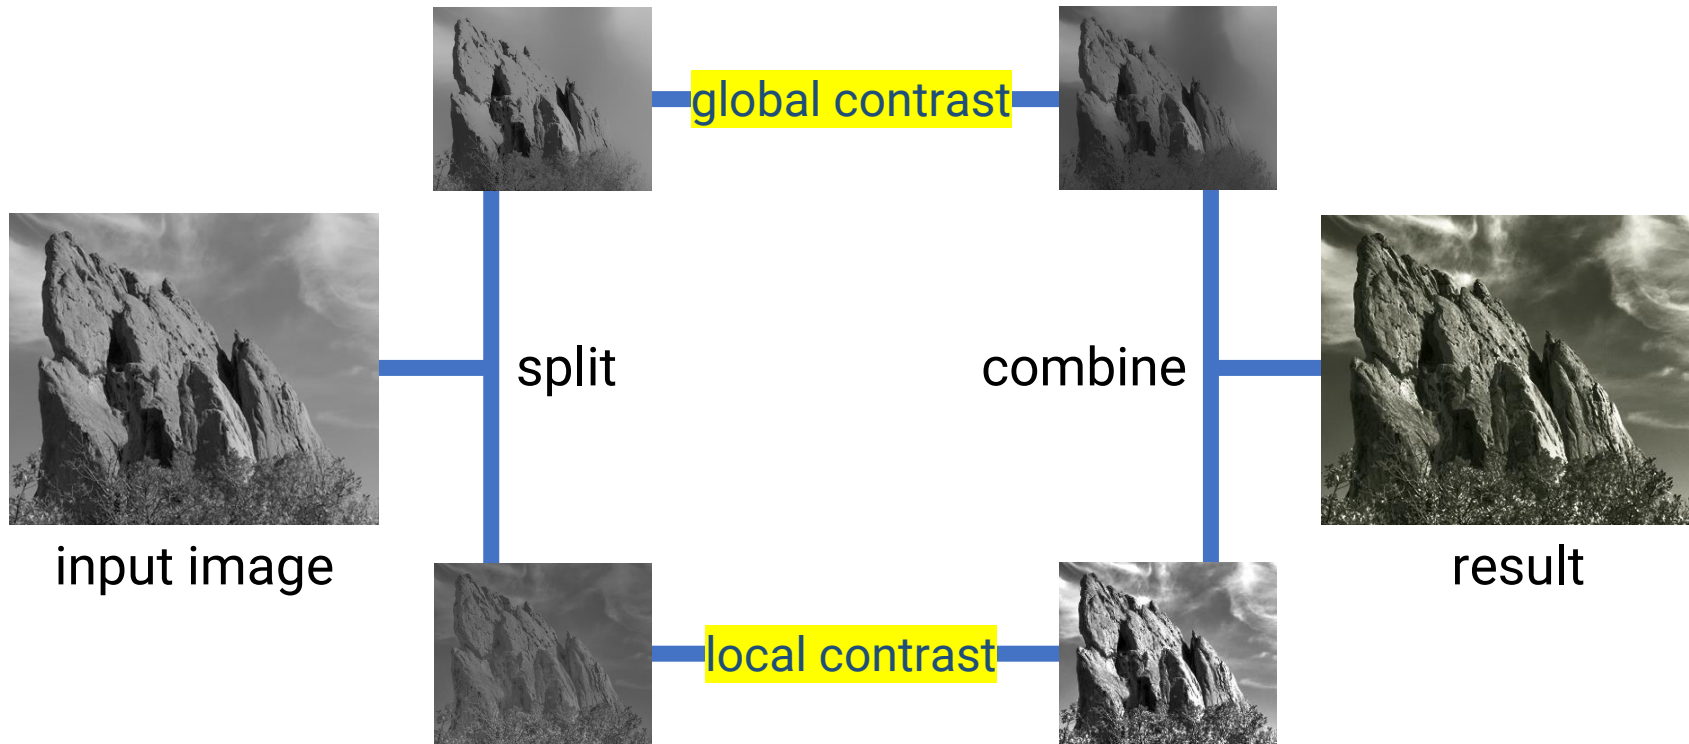
- Non-uniformly increase local contrast based on textureness



Photographic Style Transfer (cont.)

- **Algorithm**

- Combine global and local contrast




Photographic Style Transfer (cont.)

- Results



Photographic Style Transfer (cont.)

- Results




Geometrical Transformations


Types of Geometrical Transformations

- Scaling
 - Translation
 - Reflection
 - Rotation
 - Shearing
-
- For bitmapped images, we have to **transform every pixel**, and will often require the image to be **resampled**

Forward Mapping and Inverse Mapping



forward mapping



inverse mapping

Forward Mapping and Inverse Mapping

$$(x', y') = (2x, 2y)$$

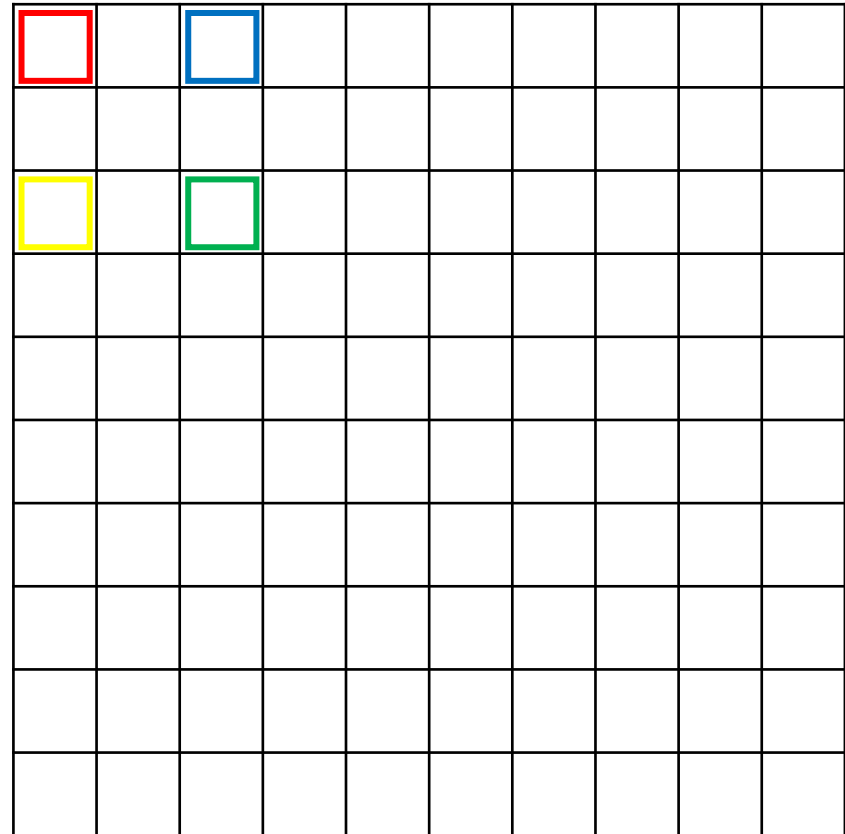



Image Scaling

- We will use image scaling as an example
- Assume we want to obtain an image which is s times larger than the original image

forward mapping $(x', y') = (sx, sy)$

inverse mapping $(x, y) = (x'/s, y'/s)$






Image Scaling (cont.)

- Three strategies to obtain an estimation of X




nearest neighbor

P_3 is closest
Use P_3 's pixel value



bilinear interpolation

$$(1-a)(1-b)P_1 + (a)(1-b)P_2 + (1-a)(b)P_3 + (a)(b)P_4$$



bicubic interpolation

using curve to
compute weight
(nonlinear)

Image Scaling (cont.)

- Example: scale an image from 160 x 120 to 800 x 600

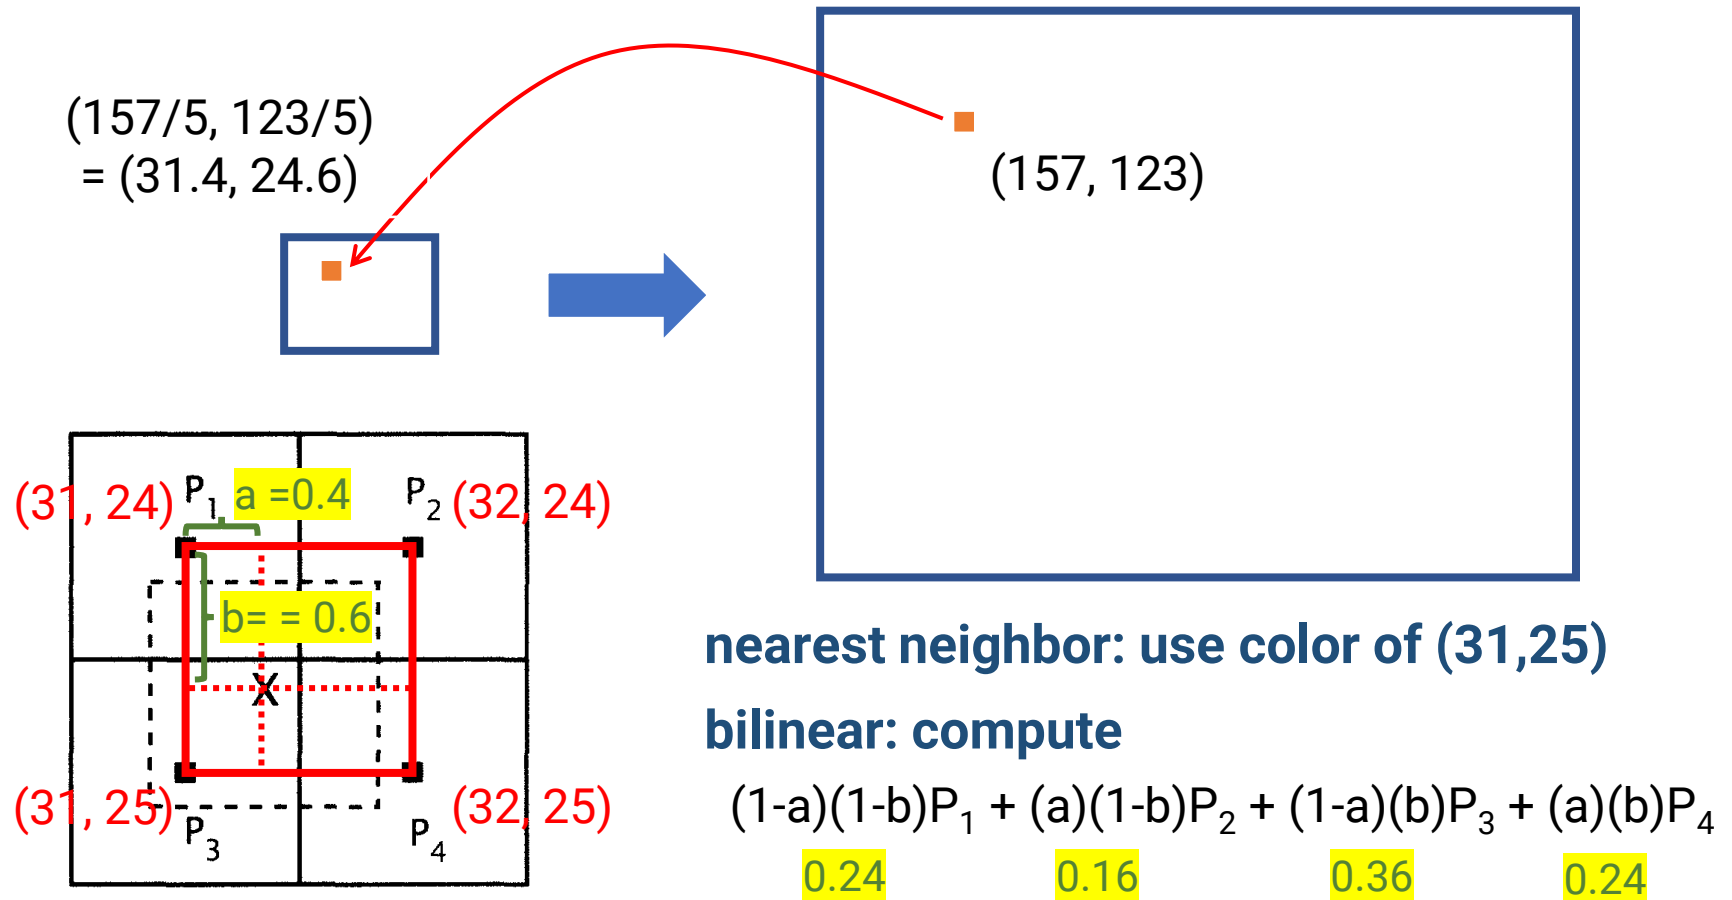




Image Scaling (cont.)

- Example



nearest neighbor



bilinear interpolation



bicubic interpolation

Image Scaling (cont.)

- Example



nearest neighbor



bilinear interpolation



bicubic interpolation

original  × 10

Image Scaling (cont.)

- Example



original



nearest neighbor



bilinear interpolation



bicubic interpolation

Image Scaling (cont.)

- OpenCV

```
void cvResize(  
    const CvArr* src,  
    const CvArr* dst,  
    int interpolation = CV_INTER_LINEAR  
);
```

interpolation can be

CV_INTER_NN, CV_INTER_LINEAR, CV_INTER_CUBIC