

#### THE UNIVERSITY OF BRITISH COLUMBIA

# **CPSC 425: Computer Vision**

#### 2 З Ο

Image Credit: https://en.wikibooks.org/wiki/Analog\_and\_Digital\_Conversion/Nyquist\_Sampling\_Rate

unless otherwise stated slides are taken or adopted from Bob Woodham, Jim Little and Fred Tung )



Lecture 7: Sampling

# Menu for Today (September 19, 2018)

#### **Topics:**

- Sampling
- Aliasing

#### **Redings:**

- Today's Lecture: [Optional] Forsyth & Ponce (2nd ed.) 4.4
- **Next** Lecture: Forsyth & Ponce (2nd ed.) 4.5

#### **Reminders:**

Assignment 1: Image Filtering and Hybrid Images due September 24th



### Bandlimited Signals Nyquist rate



### Today's "fun" Example: Clever Hans



#### Hans could get 89% of the math questions right

### Today's "fun" Example: Clever Hans



#### The course was **smart**, just not in the way van Osten thought!

#### Hans could get 89% of the math questions right

### **Assignment 1**: Image Filtering and Hybrid Images due **September 24**th

#### Common **issues** from TAs:

### 1) Don't forget to convert images to doubles!

- 2) It is convenient to also normalize the image to be between 0 and 1
- 3) Do all the operations

4) When you want to display or save convert back to uint with [0, 255] range



### Lecture 6: Re-cap

**Linear** filtering (one intepretation): new pixels are a weighted sum of original pixel values — "filter" defines weights

**Linear** filtering (another intepretation): each pixel influences the new value for itself and its neighbours — "filter" specifies the influences

**Separability** (of a 2D filter) allows for more efficient implementation (as two 1D) filters)

### Lecture 6: Re-cap

#### Convolution is associative and symmetric (correlation is not in general)

filter = boxfilter(3)signal.correlate2d(filter, filter, 'full')



#### 3x3 **Box**

1	1
1	1
1	1

=

1	2	3	2	1
2	4	6	4	2
3	6	9	6	3
2	4	6	4	2
1	2	3	2	1

3x3 **Box** 

filter = boxfilter(3)
signal.correlate2d(filter, filter, ' full')



#### 3x3 **Box**

3x3 **Box** 

1	1
1	1
1	1

=

1	2	3	2	1
2	4	6	4	2
3	6	9	6	3
2	4	6	4	2
1	2	3	2	1

Treat one filter as padded "image"



3x3 **Box** 



Treat one filter as padded "image"



3x3 **Box** 



Treat one filter as padded "image"



3x3 **Box** 



Treat one filter as padded "image"



3x3 **Box** 





#### Output

3x3 **Box** 

Treat one filter as padded "image"



3x3 **Box** 

1	1	1	
1	1	1	
1	1	1	

3x3 **Box** 

#### 

Treat one filter as padded "image"



3x3 **Box** 

# 1 1 1 1 1 1 1 1 1

3x3 **Box** 

# $=\frac{1}{81}$

1	2	3	2	1
2	4	6	4	2
3	6	9	6	3
2	4	6	4	2
1	2	3	2	1

filter = boxfilter(3)temp = signal.correlate2d(filter, filter, 'full') signal.correlate2d(filter, temp,' full')





 $\overline{256}$ 

 $\frac{1}{16}$ 

0	0	0	0	0
0	0	0	0	0
0	0	0	0	0
0	0	0	0	0
1	4	6	4	1
0	0	0	0	0
0	0	0	0	0
0	0	0	0	0
0	0	0	0	0

 $\frac{1}{16}$ 

 $\bigotimes$ 



1

 $\overline{256}$ 

1

 $\left( \right)$  $\left( \right)$ 0 0 0 0  $\left( \right)$  $\left( \right)$ 0 0 0  $\left( \right)$  $\mathbf{O}$ 1 4 1 1 6 4 16 0 0 0 0  $\left( \right)$ 0 0 0  $\left( \right)$ () $\left( \right)$  $\left( \right)$ 0  $\mathbf{O}$  $\mathbf{O}$ 

 $\frac{1}{16}$ 

 $\bigotimes$ 



 $=\frac{1}{256}$ 

1	4	6	4	1
4	16			

 $\left( \right)$ ()  $\left( \right)$  $\mathbf{0}$  $\left( \right)$  $\mathbf{O}$  $\left( \right)$  $\left( \right)$  $\left( \right)$  $\left( \right)$  $\mathbf{O}$  $\mathbf{O}$ 

 $\frac{1}{16}$ 

 $\bigotimes$ 

 $\frac{1}{256}$ 

1	4	6	4	1
4	16	24	16	4
6	24	36	24	6
4	16	24	16	4
1	4	6	4	1

 $\left( \right)$ () 0 0  $\left( \right)$ 0 0  $\mathbf{0}$ 0 0 0 0 0  $\left( \right)$  $\mathbf{O}$ 6 1 4 1 4 0 0 0 0 0 0 0 0  $\left( \right)$  $\left( \right)$ () $\left( \right)$  $\left( \right)$ 0  $\mathbf{O}$  $\mathbf{O}$ 

 $\frac{1}{16}$ 

 $\frac{1}{16}$ 

 $\bigotimes$ 

 $\frac{1}{256}$ 

1	4	6	4	1
4	16	24	16	4
6	24	36	24	6
4	16	24	16	4
1	4	6	4	1

# **Pre-Convolving** Filters

Convolving two filters of size  $m \times m$  and  $n \times n$  results in filter of size:

 $\left(n+2\left\lfloor\frac{m}{2}\right\rfloor\right)$  ×

$$\left( n+2\left\lfloor \frac{m}{2} \right\rfloor \right)$$

### **Pre-Convolving** Filters

Convolving two filters of size  $m \times m$  and  $n \times n$  results in filter of size:

$$\left(n+2\left\lfloor\frac{m}{2}\right\rfloor\right) \times \left(n+2\left\lfloor\frac{m}{2}\right\rfloor\right)$$

#### More broadly for a set of K filters of sizes $m_k \times m_k$ the resulting filter will have size:

$$\left(m_1 + 2\sum_{k=2}^{K} \left\lfloor \frac{m_k}{2} \right\rfloor\right) \times \left(m_1 + 2\sum_{k=2}^{K} \left\lfloor \frac{m_k}{2} \right\rfloor\right)$$

### Lecture 6: Re-cap

Convolution is **associative** and **symmetric** (correlation is not in general)

Convolution of a Gaussian with a Gaussian is another Gaussian

The **median filter** is a non-linear filter that selects the median in the neighbourhood

Bilateral filter is edge-preserving and non-linear. It is expressed as a product of **domain** and **range** kernels.

# Framework for Today's Topic

#### **Problem:** How do we go from the optics of image formation to digital images as arrays of numbers?

**Key Idea(s)**: Sampling and the notion of band limited functions

**Theory**: Sampling Theory

### Reminder



#### Images are a discrete, or sampled, representation of a continuous world

# What is an **Image**?

Up to now provided a physical characterization - image formation as a problem in physics/optics

Now provide a **mathematical characterization** 

- to understand how to represent images digitally
- to understand how to compute with images

### **Continuous** Case

"**Image**" suggests a 2D surface whose appearance varies from point-to-point — the surface typically is a plane (but might be curved, e.g., as is with an eye)

Appearance can be Grayscale (Black and White) or Colour

In **Grayscale**, 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

### **Continuous** Case



#### Denote the image as a function, i(x, y), where x and y are spatial variables

**Aside**: The convention for this section is to use lower case letters for the continuous case and upper case letters for the discrete case

### **Continuous** Case: Observations

-i(x,y) is a real-valued function of real spatial variables, x and y

### Recall: Pinhole Camera



Forsyth & Ponce (2nd ed.) Figure 1.2

### **Continuous** Case: 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  $0 \le i(x,y) \le M$ 

for some maximum brightness  ${\cal M}$ 

### **Continuous** Case: 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

for some maximum brightness M

-i(x,y) is **bounded in extent**. That is, i(x,y) is non-zero (i.e., strictly positive) over, at most, a bounded region

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

## **Continuous** Case

where x and y are spatial variable and t is a **temporal variable** 

- To make the dependence of brightness on wavelength explicit, we can instead write  $i(x, y, t, \lambda)$  where x, y and t are as above and where  $\lambda$  is a spectral variable

More commonly, we think of "color" already as discrete and write

for specific colour channels, R, G and B

# - Images also can be considered a function of time. Then, we write i(x, y, t)

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

34

Idea: Superimpose (regular) grid on continuous image



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



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



Denote the discrete image as I(X, Y)

We can store the pixels in a matrix or array

#### **Question**: How to sample?

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

#### **Answer**:

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

**Question**: What about the brightness samples themselves?

**Question**: What about the brightness samples themselves?

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

Recall: 
$$0 \le i(x, y) \le M$$

We divide the range [0, M] into a finite number of equivalence classes. This is called **quantization**.

The values are called **grey-levels**.

Quantization is a topic in its own right

- For now, a simple linear scheme is sufficient
- evenly spaced intervals as follows:

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

where  $\lfloor$  is floor (i.e., greatest integer less than or equal to) Typically n = 8 resulting in grey-levels in the range [0, 255]

Suppose n bits-per-pixel are available. One can divide the range [0, M] into

0	Đ	0	0	0	0	0
o	ð	0	D	0	0	0
ä	Ô	٥	٥	C	Ċ	Q
¢	¢	¢	¢	¢	0	¢
σ	¢	Ō	D	0	0	O
o	Ð	o	0	0	0	0
0	Ð	Ø	0	0	0	0
٩	¢	¢	٥	0	0	٥
_				~		
¢				0		
¢				0		
¢				0		
¢				0		
0				0		
¢				0		
0				0		
¢				¢		

#### It is clear that some information may be lost when we work on a discrete pixel grid.



#### Forsyth & Ponce (2nd ed.) Figure 4.7 42



0	Ð	Ð	0	0	0	0
0	ð	0	D	0	0	a
a	٥	Ō	Ď	Ċ	Ċ	a
a	¢	¢	٥	¢	٥	¢
σ	¢	Ō	D	0	0	Q
0	Ð	o	0	0	0	0
0	0	o	0	0	0	0
٩	٥	٥	٥	0	0	٥
Q				0		



#### It is clear that some information may be lost when we work on a discrete pixel grid.

#### Forsyth & Ponce (2nd ed.) Figure 4.7 43



0	Đ	0	0	0	0	0
o	ð	0	D	0	0	0
ä	Ô	٥	٥	C	Ċ	Q
à	¢	¢	Φ	¢	0	¢
σ	¢	Ō	D	0	0	O
o	Ð	o	0	0	0	0
0	Ð	Ø	0	0	0	0
٩	¢	¢	٥	0	0	٥
_				~		
¢				0		
¢				0		
¢				0		
¢				0		
0				0		
¢				0		
0				0		
¢				¢		

#### It is clear that some information may be lost when we work on a discrete pixel grid.





#### Forsyth & Ponce (2nd ed.) Figure 4.7 44



0	Đ	0	0	0	0	0
o	ð	0	D	0	0	0
ä	Ô	٥	٥	C	Ċ	Q
à	¢	¢	Φ	¢	0	¢
σ	¢	Ō	D	0	0	O
o	Ð	o	0	0	0	0
0	Ð	Ø	0	0	0	0
٩	¢	¢	٥	0	0	٥
_				~		
¢				0		
¢				0		
¢				0		
¢				0		
0				0		
¢				0		
¢				0		
¢				¢		

Forsyth & Ponce (2nd ed.) Figure 4.7 45

#### It is clear that some information may be lost when we work on a discrete pixel grid.







0	Ð	Ð	0	0	0	0
0	¢	0	o	0	0	a -
a	0	ð	٥	Ď	a	٥
¢	¢	¢	¢	¢	0	¢
σ	Q	0	D	0	0	a
0	o	0	0	0	0	0
0	0	o	0	0	0	0
a	0	ð	D	0	0	٥
¢				0		
¢				¢		
¢				0		

Forsyth & Ponce (2nd ed.) Figure 4.7 46

#### It is clear that some information may be lost when we work on a discrete pixel grid.







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

**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)?

**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** 

**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

**Case 0**: Suppose i(x, y) = k (with k being one of our gray levels)



Note: we use equidistant sampling at integer values for convenience, in general, sampling doesn't need to be equidistant

**Case 0**: Suppose i(x, y) = k (with k being one of our gray levels)



X

**Case 0**: Suppose i(x, y) = k (with k being one of our gray levels)



I(X,Y) = k. Any standard interpolation function would give i(x,y) = k for noninteger x and y (irrespective on how coarse the sampling is)



#### **Case 0**: Suppose i(x, y) has a discontinuity not falling precisely at integer x, y



We cannot reconstruct i(x, y) exactly because we can never know exactly where the discontinuity lies

#### **Case 0**: Suppose i(x, y) has a discontinuity not falling precisely at integer x, y



We cannot reconstruct i(x, y) exactly because we can never know exactly where the discontinuity lies

#### **Case 0**: Suppose i(x, y) has a discontinuity not falling precisely at integer x, y

### This is **impossible**!

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

Next, we build intuition based on informal argument

- between samples
- "rate of change" means derivative
- the formal concept is **bandlimited signal**
- "bandlimit" and "constraint on derivative" are linked
- Think of music
- bandlimited if it has some maximum temporal frequency
- the upper limit of human hearing is about 20 kHz
- Think of imaging systems. Resolving power is measured in
- "line pairs per mm" (for a bar test pattern)
- "cycles per mm" (for a sine wave test pattern)
- An image is bandlimited if it has some maximum spatial frequency

#### Exact reconstruction requires constraint on the rate at which i(x,y) can change

58

How do we discretize the signal?





How do we discretize the signal?



How do we discretize the signal?



How many samples should I take? Can I take as many samples as I want?

How do we discretize the signal?



How many samples should I take? Can I take as few samples as I want?

How do we discretize the signal?



#### Signal can be confused with one at lower frequency

How do we discretize the signal?



Signal can be confused with one at lower frequency

How do we discretize the signal?



#### Signal can always be confused with one at higher frequency

# Undersampling = Aliasing



The challenge to intuition is the fact that music (in the 1D case) and images (in the 2D case) can be represented as linear combinations of individual sine waves of differing frequencies and phases (remember discussion on FFTs)

A fundamental result (**Sampling Theorem**) is: For bandlimited signals, if you sample regularly at or above twice the maximum frequency (called the Nyquist rate), then you can reconstruct the original signal exactly