Data Sience: Introduction to Image Processing

Bruno Galerne

Monday January 8, 2024

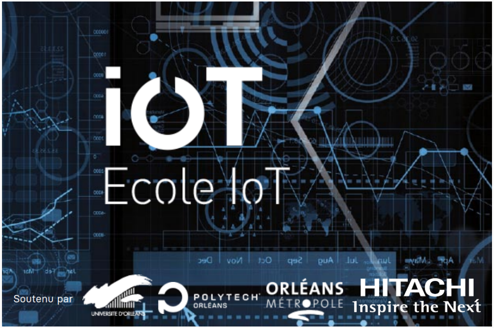

Courses

Content: 6 courses

3 courses on image processing:

1. Introduction to image processing
2. Histogram and spatial filtering
3. Spectral filtering (Fourier) and introduction to variational formulation

3 courses on neural networks:

1. Introduction to neural networks
2. Convolutional neural networks for image classification
3. Deep neural networks for computer vision

Disclaimer

Several slides from **Charles Deledalle's** course "UCSD ECE285 Image and video restoration" (30 $\times$ 50 minutes course) given at UCSD (University of California, San Diego).

`www.charles-deledalle.fr/`

This is an advanced course to go further into the mathematics of image processing.

Introduction: What is image processing?

- Imaging:

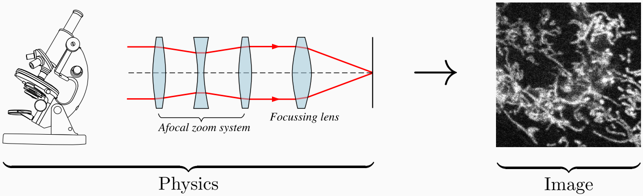

Modeling the image formation process

Imaging sciences – Overview

- Imaging:

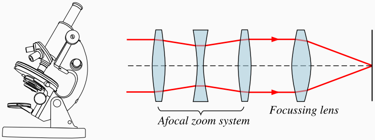

→

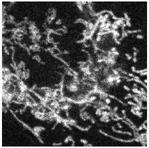

Physics

Image

Modeling the image formation process

- Computer graphics:

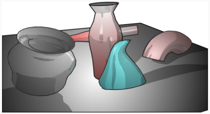

→

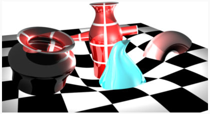

Symbols

Image

Rendering images/videos from symbolic representation

- Computer vision:

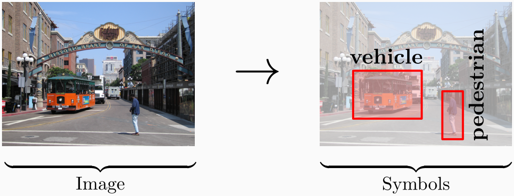

Extracting information from images/videos

Imaging sciences – Overview

- Computer vision:

→

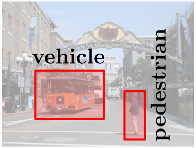

Image

Symbols

Extracting information from images/videos

- Image/Video processing:

→

Image

Image

Producing new images/videos from input images/videos

Imaging sciences – Image processing

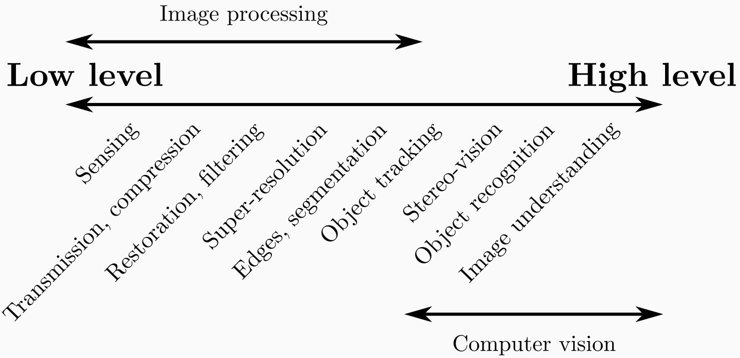

Imaging sciences – Image processing

Denoising

Enhancement

Compression

Feature detection

Inpainting

Super-resolution

Source: Iasonas Kokkinos

- Image processing: define a new image from an existing one
- Video processing: same problems + motion information

Denoising

Enhancement

Compression

Feature detection

Inpainting

Super-resolution

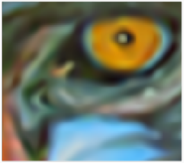

Source: Iasonas Kokkinos

- Image processing: define a new image from an existing one
- Video processing: same problems + motion information

Geometric transform

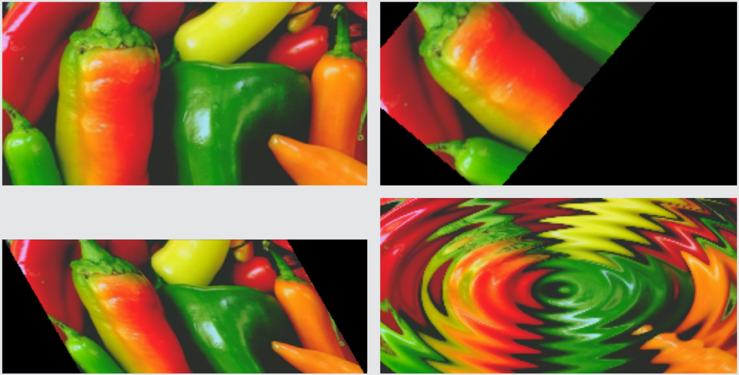

Change pixel location

Colorimetric transform

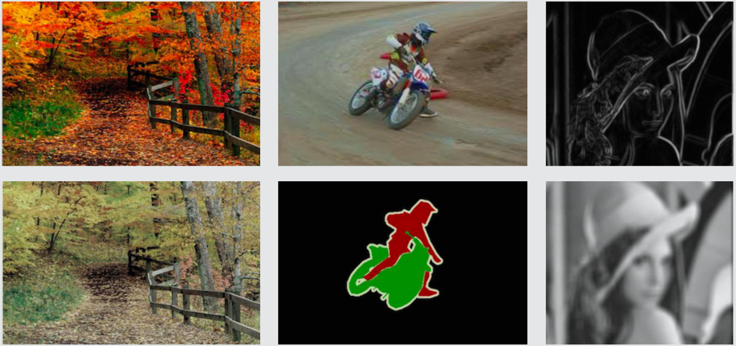

- Filtering: change pixel values
- Segmentation: provide an attribute to each pixel

Imaging sciences – Photo manipulation

Photo manipulation – Applications & Techniques (sources Wikipedia)

Media industry

Skin flaw removal (Minnie Driver by Justin Hoch)

Art

Editing (by Achraf Baznani)

Propaganda

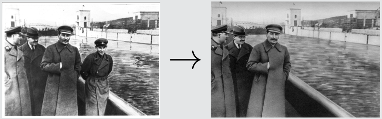

Joseph Stalin with Nikolai Yezhov entirely removed after retouching

- Media / Journalism / Advertising
- Restoration of cultural heritage
- Propaganda / Political purpose
- Art / Personal use

Photo manipulation – Applications & Techniques

(sources Wikipedia)

Media industry

→

Skin flaw removal (Minnie Driver by Justin Hoch)

Art

Editing (by Achraf Baznani)

Propaganda

→

Joseph Stalin with Nikolai Yezhov entirely removed after retouching

- Media / Journalism / Advertising
- Restoration of cultural heritage
- Propaganda / Political purpose
- Art / Personal use

- Color & contrast enhancement
- Image sharpening (reduce blur)
- Removing elements (inpainting)
- Removing flaws (skin, scratches)
- Image compositing/fusion
- Image colorization

Photo manipulation – Applications & Techniques

(sources Wikipedia)

Media industry

Skin flaw removal (Minnie Driver by Justin Hoch)

Art

Editing (by Achraf Baznani)

Propaganda

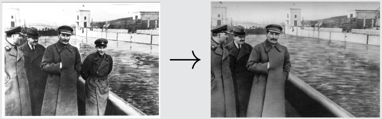

Joseph Stalin with Nikolai Yezhov entirely removed after retouching

- Media / Journalism / Advertising
- Restoration of cultural heritage
- Propaganda / Political purpose
- Art / Personal use

- Color & contrast enhancement
- Image sharpening (reduce blur)
- Removing elements (inpainting)
- Removing flaws (skin, scratches)
- Image compositing/fusion
- Image colorization

Often handmade by graphic designers/artists/confirmed amateurs
or aided with raster images/graphics editor

Classical editors: Adobe Photoshop (commercial), GIMP (free and open-source)

Imaging sciences – Is image processing = Photo manipulation?

Photo manipulation

- Manual/Computer aided
- Performed image per image
- Users: artists, graphic designers
- Target: general public
- Input: photography
- Goal: visual aspects

v.s

Main image processing purposes

- Automatic/Semi-supervised
- Applied to image datasets
- Users: industry, scientists
- Target: industry, sciences
- Input: any kind of $\geqslant$ 2d signals
- Goal: measures, post analysis

v.s

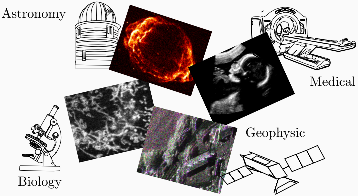

Photo manipulation uses some image processing tools

Scope of image processing is much wider than photography

Multidisciplinary of Image processing

Intersection of several covering fields

- **Physics and biology:** link between phenomena and measures
- **Mathematics:** analyze observations and make predictions
- **Computer science:** algorithms to extract information
- **Statistics:** account for uncertainties in data

Imaging sciences – Related fields

Multidisciplinary of Image processing

Intersection of several covering fields

- **Physics and biology:** link between phenomena and measures
- **Mathematics:** analyze observations and make predictions
- **Computer science:** algorithms to extract information
- **Statistics:** account for uncertainties in data

Differences with signal processing

- Image processing: subset of signal processing
- Inputs and outputs: images, series of images or videos
- Content: sound waves, stock prices behave differently
- Signals are usually causal: $f(t_0)$ depends only on $f(t)$ for any time $t \leqslant t_0$
- Images are non-causal: $f(s_0)$ may depend on $f(s)$ for any position s

Imaging sciences – What is image restoration?

What is image restoration?

- Subset of image processing
- Input: corrupted image
- Output: estimate of the clean/original image
- Goal: reverse the degradation process

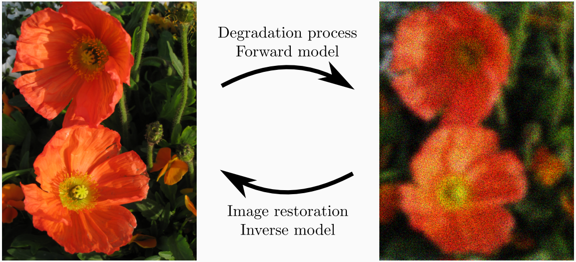

Image restoration requires **accurate models** for the degradation process.
Knowing and modeling the sources of corruptions is essential.

Why image restoration?

- Artistic value?
- or, Automatic image analysis?
  - Object recognition
  - Image indexation
  - Image classification
  - ...
- Usually one of the first steps in computer vision (CV) pipelines.

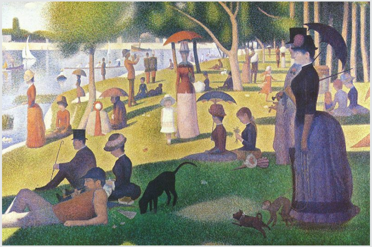

Pointillism (Georges Seurat, 1884-1886)

- A source of inspiration to perform higher level tasks.

What is an image?

La Trahison des images, René Magritte, 1928
(Los Angeles County Museum of Art)

Imaging sciences – What is an image for us?

A function?

- Think of an image as a function f from $\mathbb{R}^2$ (2d space) to $\mathbb{R}$ (values).
- $f(s_1, s_2)$ gives the intensity at location $(s_1, s_2) \in \mathbb{R}^2$.
- In practice, usually limited to: $f : [0, 1]^2 \to \mathbb{R}$.

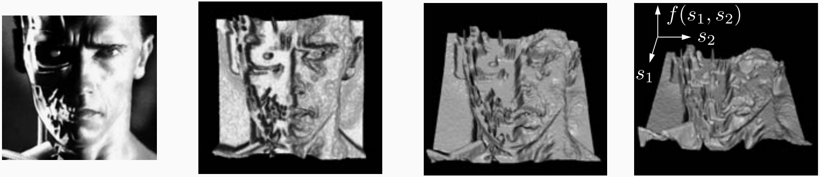

Source: Steven Seitz

Convention: larger values correspond to brighter colors.

Imaging sciences – What is an image for us?

A function?

- Think of an image as a function f from $\mathbb{R}^2$ (2d space) to $\mathbb{R}$ (values).
- $f(s_1, s_2)$ gives the intensity at location $(s_1, s_2) \in \mathbb{R}^2$.
- In practice, usually limited to: $f : [0,1]^2 \to \mathbb{R}$.

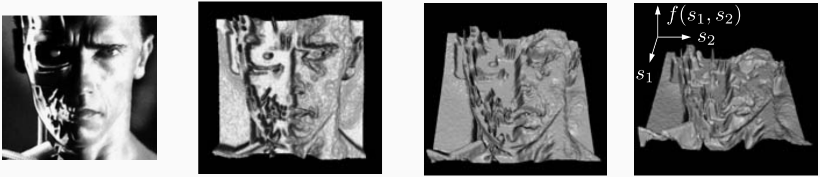

Source: Steven Seitz

Convention: larger values correspond to brighter colors.

A color image is defined similarly as a 3 component vector-valued function:

$$f(s_1, s_2) = \begin{pmatrix} r(s_1, s_2) \\ g(s_1, s_2) \\ b(s_1, s_2) \end{pmatrix} .$$

Imaging sciences – Types of images

- Continuous images:
  - Analog images/videos,
  - Vector graphics editor, or (Adobe Illustrator, Inkscape, . . .)
  - 2d/3d+time graphics editors. (Blender, 3d Studio Max, . . .)
  - Format: `svg`, `pdf`, `eps`, `3ds`. . .
- Discrete images:
  - Digital images/videos,
  - Raster graphics editor. (Adobe Photoshop, GIMP, . . .)
  - Format: `jpeg`, `png`, `ppm`. . .
- All are displayed on a digital screen as <u>a digital image/video</u> (rendering).

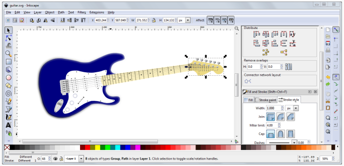

(a) Inkscape

(b) Gimp

- Progressively changing recording medium,
- Often chemical or electronic,
- Modeled as a continuous signal, *e.g.*:
  - Gray level images: $[0,1]^2 \to \mathbb{R}$ position to gray level,
  - Color images: $[0,1]^2 \to \mathbb{R}^3$ position to RGB levels.

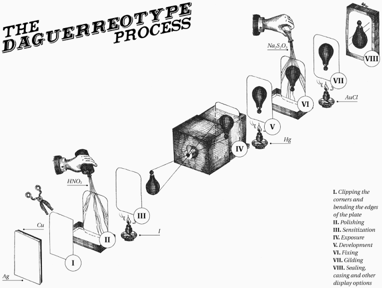

(a) Daguerrotype

(b) Roll film

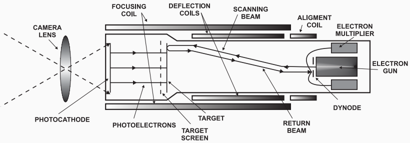

(c) Orthicon tube

Example (Analog photography/video)

- First type of photography was analog.

(a) Daguerrotype

(b) Carbon print

(c) Silver halide

- Still in used by photographs and the movie industry for its artistic value.

(d) Carol (2015, Super 16mm)

(e) Hateful Eight (2015, 70mm)

(f) Grand Budapest Hotel (2014, 35mm)

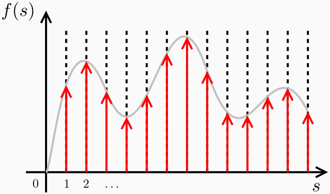

Raster images

- Sampling: reduce the 2d continuous space to a discrete grid $\Omega \subseteq \mathbb{Z}^2$
- Gray level image: $\Omega \to \mathbb{R}$ (discrete position to gray level)
- Color image: $\Omega \to \mathbb{R}^3$ (discrete position to RGB)

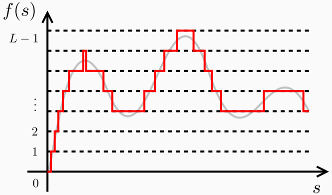

Bitmap image

- Quantization: map each value to a discrete set $[0, L-1]$ of L values (*e.g.*, round to nearest integer)

- Often $L = 2^8 = 256$ (8bit images $\equiv$ `unsigned char`)
  - Gray level image: $\Omega \to [0, 255]$ $(255 = 2^8 - 1)$
  - Color image: $\Omega \to [0, 255]^3$

- Optional: assign instead an index to each pixel pointing to a color palette (format: `.png`, `.bmp`)

Digital imagery

- Digital images: sampling + quantization:

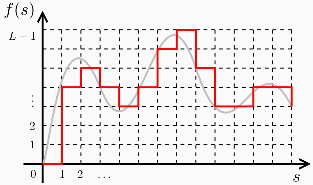

$\longrightarrow$ 8bit images can be seen as a matrix of integer values

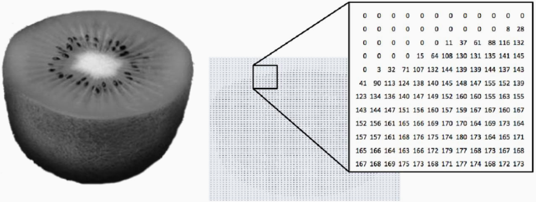

We will refer to an element $s \in \Omega$ as a pixel location, $x(s)$ as a pixel value, and the pair $(s, x(s))$ as a pixel ("picture element").

Functional representation: $f : \Omega \subseteq \mathbb{Z}^d \to \mathbb{R}^K$

- d: dimension ($d = 2$ for pictures, $d = 3$ for videos, …)
- K: number of channels ($K = 1$ monochrome, 3 color, …)
- $s = (i, j)$: pixel position in Ω
- $f(s) = f(i, j)$: pixel value(s) in $\mathbb{R}^K$

Functional representation: $f : \Omega \subseteq \mathbb{Z}^d \to \mathbb{R}^K$

- d: dimension ($d = 2$ for pictures, $d = 3$ for videos, ...)
- K: number of channels ($K = 1$ monochrome, 3 color, ...)
- $s = (i, j)$: pixel position in Ω
- $f(s) = f(i, j)$: pixel value(s) in $\mathbb{R}^K$

Array representation ($d = 2$): $x \in (\mathbb{R}^K)^{n_1 \times n_2}$

- $n_1 \times n_2$: n_1: image height, and n_2: width
- $x_{i,j} \in \mathbb{R}^K$: pixel value(s) at position $s = (i, j)$: $x_{i,j} = f(i, j)$

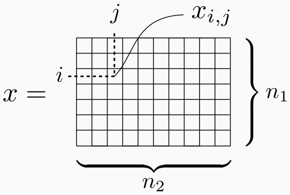

Vector representation: $y \in (\mathbb{R}^K)^n$

- $n = n_1 \times n_2$: image size (number of pixels)
- $y_k \in \mathbb{R}^K$: value(s) of the k-th pixel at position s_k: $y_k = f(s_k)$

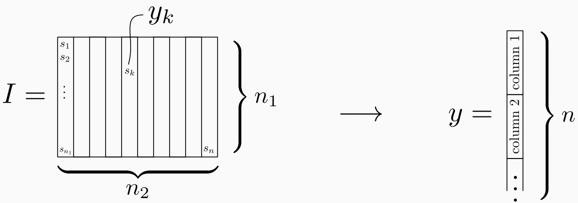

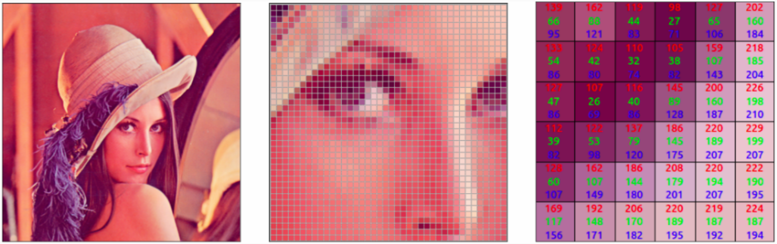

Color 2d image: $\Omega \subseteq \mathbb{Z}^2 \to [0, 255]^3$

- Red, Green, Blue (RGB), $K = 3$
- RGB: Usual colorspace for acquisition and display
- There exist other colorspaces for different purposes:
  HSV (Hue, Saturation, Value), YUV, YCbCr. . .

Spectral image: $\Omega \subseteq \mathbb{Z}^2 \to \mathbb{R}^K$

- Each of the K channels is a wavelength band
- For $K \approx 10$: multi-spectral imagery
- For $K \approx 200$: hyper-spectral imagery
- Used in astronomy, surveillance, mineralogy, agriculture, chemistry

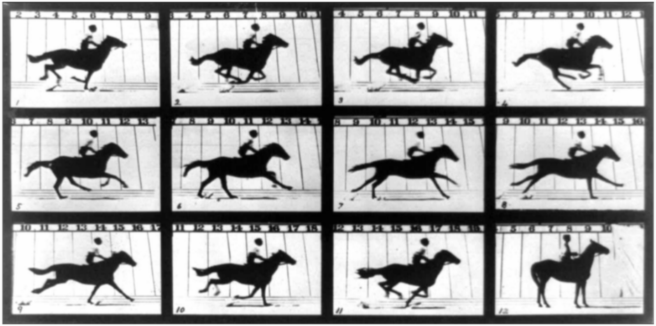

The Horse in Motion (1878, Eadweard Muybridge)

Gray level video: $\Omega \subseteq \mathbb{Z}^3 \to \mathbb{R}$

- 2 dimensions for space
- 1 dimension for time

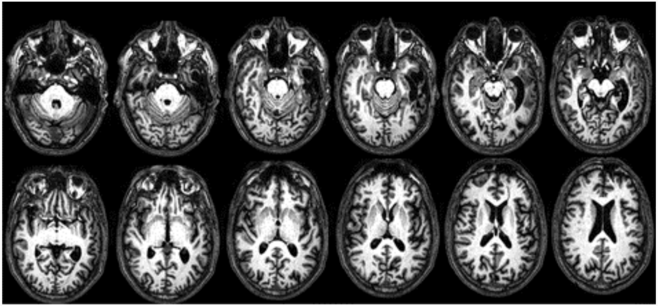

MRI slices at different depths

3d brain scan: $\Omega \subseteq \mathbb{Z}^3 \to \mathbb{R}$

- 3 dimensions for space
- 3d pixels are called voxels ("volume elements")

What is noise?

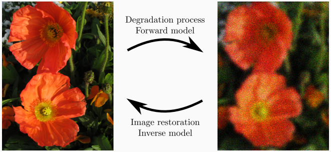

Knowing and modeling the sources of corruptions is essential.

Analog optical imagery

Basic principle of silver-halide photography

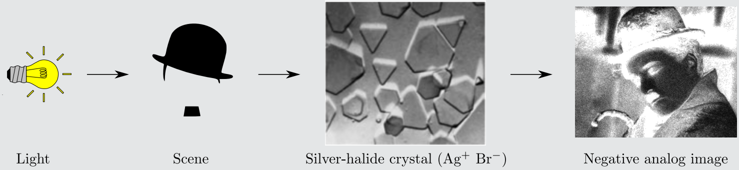

Crystals are sensitive to light
(chemical reaction during exposure and development)

Film grain:

- Depends on the amount of crystals (quality/type of film roll)
- Depends on the scale it is observed (noticeable in an over-enlarged picture)

Analog optical imagery

Analog television

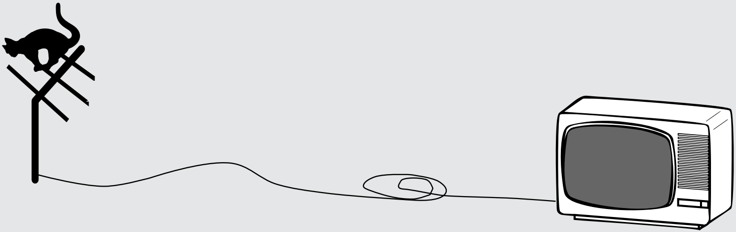

Noise due to bad transmission and/or interference

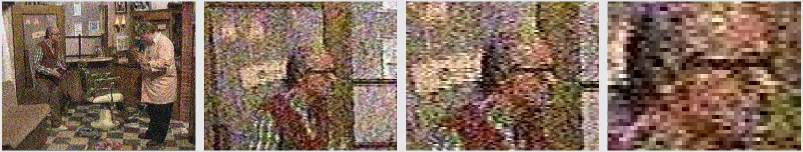

Digital optical imagery / CCD

Include:

- digital photography
- optical microscopy
- optical telescopes (e.g., Hubble, Planck, . . .)
- optical earth observation satellite (e.g., Landsat, Quickbird, . . .)

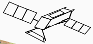

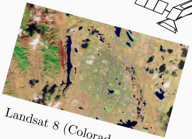

Landsat 8 (Colorado, USA)

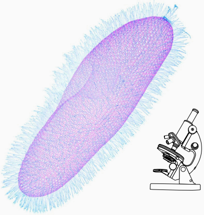

Leica microscope (paramecium aurelia)

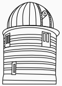

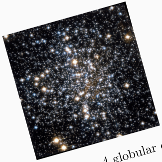

Hubble (Messier 4 globular cluster)

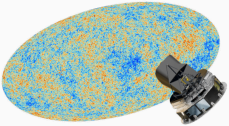

Planck (cosmic microwave background)

Digital optical imagery / CCD

Charge Coupled Device – Simplified description

 Some photons,

 captured during the exposure time (shutter speed),

 are converted to electrons,

 leading to a charge converted to voltage,

 next amplified,

 quantized and digitized,

providing a grey level.

Charge Coupled Device – Simplified description

- Some photons,
- captured during the exposure time (shutter speed), ... Shot noise
- are converted to electrons, ... Thermal noise
- leading to a charge converted to voltage, ... Readout noise
- next amplified, ... ISO sensitivity
- quantized and digitized, ... ISO sensitivity
- providing a grey level. (Often followed by non-linear post-processing and lossy compression)

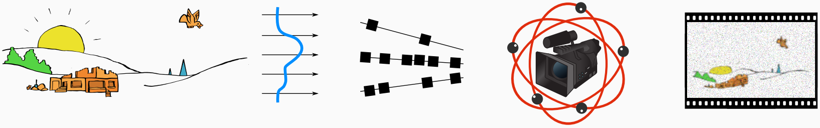

Scene — Light intensity — Photon emission — Electronic fluctuations — Digital Image

Random fluctuations lead to noise

Digital optical imagery – Noise modeling

- Take several pictures of the same scene, and focus on one given pixel,
- There are always unwanted fluctuations around the “true” pixel value,
- These fluctuations are called noise,
- Usually described by a probability density or mass function (pdf/pmf),
- Stochastic process Y parametrized by a deterministic signal of interest x.

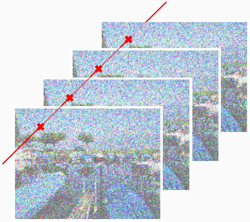

- Take several pictures of the same scene, and focus on one given pixel,
- There are always unwanted fluctuations around the “true” pixel value,
- These fluctuations are called noise,
- Usually described by a probability density or mass function (pdf/pmf),
- Stochastic process Y parametrized by a deterministic signal of interest x.

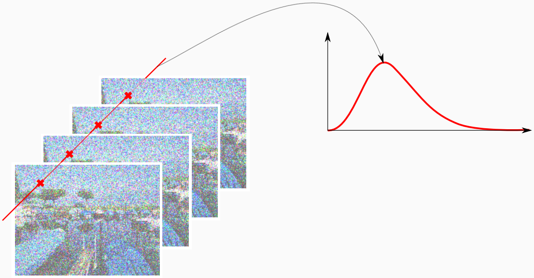

Digital optical imagery – Noise modeling

- Take several pictures of the same scene, and focus on one given pixel,
- There are always unwanted fluctuations around the "true" pixel value,
- These fluctuations are called noise,
- Usually described by a probability density or mass function (pdf/pmf),
- Stochastic process Y parametrized by a deterministic signal of interest x.

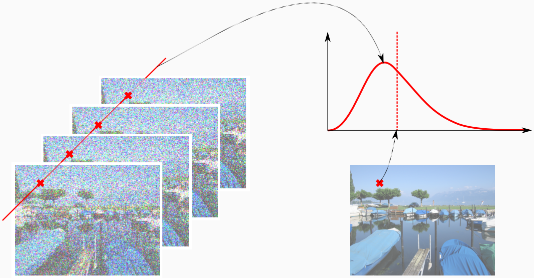

Digital optical imagery – Noise modeling

- Take several pictures of the same scene, and focus on one given pixel,
- There are always unwanted fluctuations around the “true” pixel value,
- These fluctuations are called noise,
- Usually described by a probability density or mass function (pdf/pmf),
- Stochastic process Y parametrized by a deterministic signal of interest x.

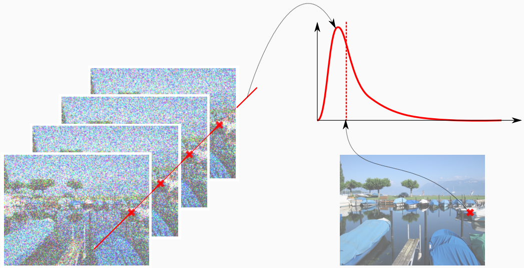

Digital optical imagery – Noise modeling

- Take several pictures of the same scene, and focus on one given pixel,
- There are always unwanted fluctuations around the "true" pixel value,
- These fluctuations are called noise,
- Usually described by a probability density or mass function (pdf/pmf),
- Stochastic process Y parametrized by a deterministic signal of interest x.

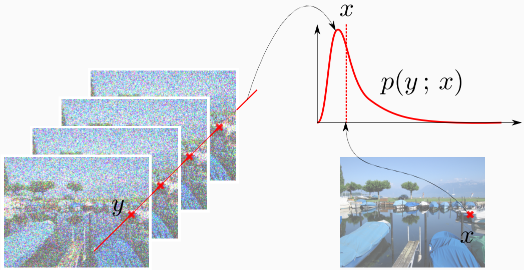

x true unknown pixel value, y noisy observed value (a realization of Y),
link: $p_Y(y\,;\,x)$ noise model

Shot noise

- Number of captured photons $y \in \mathbb{N}$ fluctuates around the signal of interest

$$x = PQ_et$$

  - x: expected quantity of light
  - Q_e: quantum efficiency (depends on wavelength)
  - P: photon flux (depends on light intensity and pixel size)
  - t: integration time
- Variations depend on exposure times and light conditions.

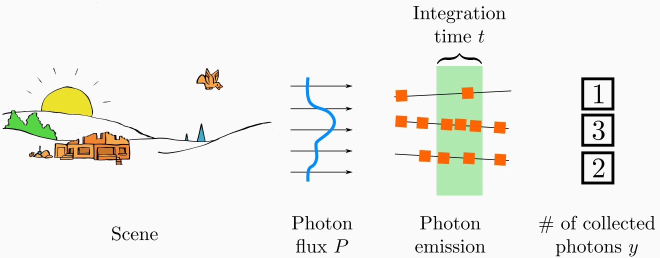

Shot noise

- Number of captured photons $y \in \mathbb{N}$ fluctuates around the signal of interest

$$x = P Q_e t$$

  - x: expected quantity of light
  - Q_e: quantum efficiency (depends on wavelength)
  - P: photon flux (depends on light intensity and pixel size)
  - t: integration time

- Variations depend on exposure times and light conditions.

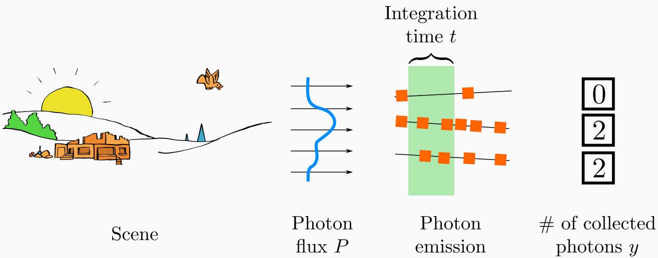

Digital optical imagery – Shot noise

Shot noise and Poisson distribution

- Distribution of Y modeled by the Poisson distribution

$$p_Y(y\,;\,x) = \frac{x^y e^{-x}}{y!}$$

- Number of photons $y \in \mathbb{N}$ fluctuates around the signal of interest $x \in \mathbb{R}$

$$\mathbb{E}[Y] = \sum_{y=0}^{\infty} y p_Y(y\,;\,x) = x$$

- Fluctuations proportional to $\mathrm{Std}[Y] = \sqrt{\mathrm{Var}[Y]} = \sqrt{x}$

$$\mathrm{Var}[Y] = \sum_{y=0}^{\infty} (y - x)^2 p_Y(y\,;\,x) = x$$

- Inherent when counting particles in a given time window

We write $Y \sim \mathcal{P}(x)$

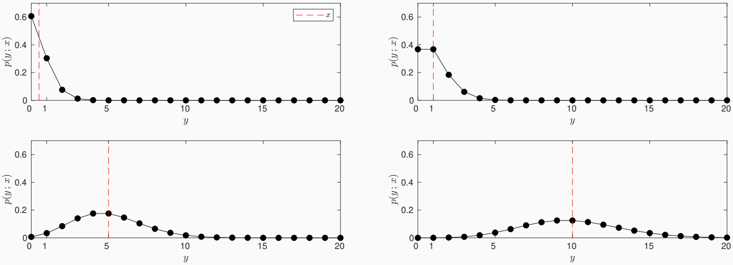

Figure 1 – Distribution of Y for a given quantity of light x

- For $x = 0.5$: mostly 0 photons, Spread ≈ 0.7
- For $x = 1$: mostly 0 or 1 photons, Spread $= 1$
- For $x \gg 1$: bell shape around x, Spread $= \sqrt{x}$

Spread is higher when $x = PQ_e t$ is large.
Should we prefer small exposure time t? and lower light conditions P?

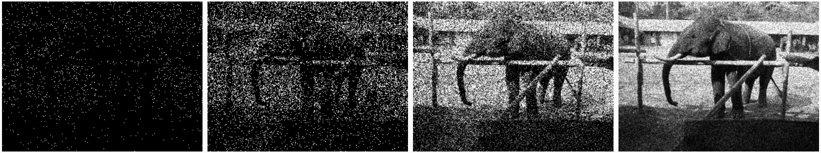

(a) Peak $= 0.05$ (b) Peak $= 0.40$ (c) Peak $= 3.14$ (d) Peak $= 24.37$

Figure 2 – Aspect of shot noise under different light conditions. Peak $= \max_i x_i$.

Digital optical imagery – Shot noise

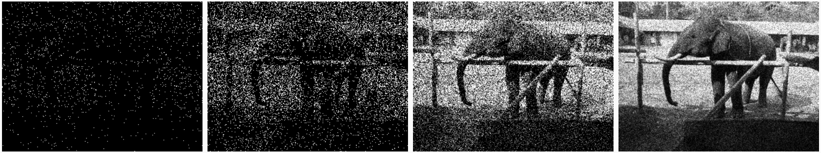

(a) Peak $= 0.05$ (b) Peak $= 0.40$ (c) Peak $= 3.14$ (d) Peak $= 24.37$

Figure 2 – Aspect of shot noise under different light conditions. Peak $= \max_i x_i$.

Signal to Noise Ratio

$$\mathrm{SNR} = \frac{x}{\sqrt{\mathrm{Var}[Y]}}, \quad \text{for shot noise} \quad \mathrm{SNR} = \sqrt{x}$$

- Measure of difficulty/quality
- The higher the easier/better
- Rose criterion: an SNR of at least 5 is needed to be able to distinguish image features at 100% certainty.

The spread (variance) is not informative,
what matters is the spread relatively to the signal (SNR)

Readout noise (a.k.a, electronic noise)

- Inherent to the process of converting CCD charges into voltage
- Measures $y \in \mathbb{R}$ fluctuate around a voltage $x \in \mathbb{R}$

$$\mathbb{E}[Y] = \int y p_Y(y\,;\,x)\,\mathrm{d}y = x$$

- Fluctuations are independent of x

$$\mathrm{Var}[Y] = \int (y - x)^2 p_Y(y\,;\,x)\,\mathrm{d}y = \sigma^2$$

- Described as Gaussian distributed

$$p_Y(y\,;\,x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(y-x)^2}{2\sigma^2}\right)$$

- Additive behavior: $Y = x + W, \quad W \sim \mathcal{N}(0, \sigma^2)$

We write $Y \sim \mathcal{N}(x, \sigma^2)$

Gaussian/Normal distribution

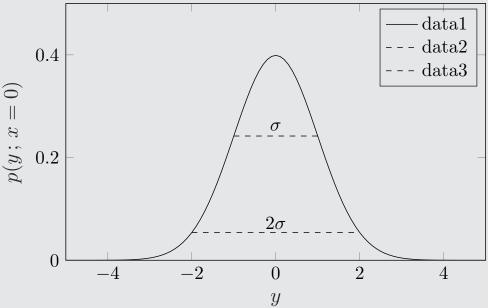

- Symmetric with bell shape.
- Common to models $\pm\sigma$ uncertainties with very few outliers
  $\mathbb{P}[|Y-x| \leqslant \sigma] \approx 0.68$, $\mathbb{P}[|Y-x| \leqslant 2\sigma] \approx 0.95$, $\mathbb{P}[|Y-x| \leqslant 3\sigma] \approx 0.99$.
- Arises in many problems due to the Central Limit Theorem.
- Simple to manipulate: eases computation in many cases.

Digital optical imagery – Shot noise vs Readout noise

Shot noise is signal-dependent (Poisson noise)

Readout noise is signal-independent (Gaussian noise)

Digital optical imagery – Thermal and total noise

Thermal noise (a.k.a, dark noise)

- Number of generated electrons fluctuates with the CCD temperature
- Additive Poisson distributed: $Y = x + N \quad \text{with} \quad N \sim \mathcal{P}(\lambda)$
- Signal independent

Digital optical imagery – Thermal and total noise

Thermal noise (a.k.a, dark noise)

- Number of generated electrons fluctuates with the CCD temperature
- Additive Poisson distributed: $Y = x + N \quad \text{with} \quad N \sim \mathcal{P}(\lambda)$
- Signal independent

Total noise in CCD models

$$Y = Z + N + W$$

$$\text{with} \quad \begin{cases} Z \sim \mathcal{P}(x), \\ N \sim \mathcal{P}(\lambda), \\ W \sim \mathcal{N}(0, \sigma^2). \end{cases}$$

$$\text{SNR} = \frac{x}{\sqrt{x + \lambda + \sigma^2}}$$

where $\quad x = PQ_e t, \quad \lambda = Dt$

- t: exposure time
- P: photon flux per pixel *(depends on luminosity)*
- Q_e: quantum efficiency *(depends on wavelength)*
- D: dark current *(depends on temperature)*
- σ: readout noise *(depends on electronic design)*

Digital optical imagery – How to reduce noise?

$$\text{SNR} = \frac{x}{\sqrt{x + \lambda + \sigma^2}} \quad \text{where} \quad x = PQ_e t, \quad \lambda = Dt$$

Photon noise

- Cannot be reduced via camera design
- Reduced by using a longer exposure time t
- Reduced by increasing the scene luminosity, higher P (e.g., using a flash)
- Reduced by increasing the aperture, higher P

Digital optical imagery – How to reduce noise?

$$\text{SNR} = \frac{x}{\sqrt{x + \lambda + \sigma^2}} \quad \text{where} \quad x = PQ_e t, \quad \lambda = Dt$$

Photon noise

- Cannot be reduced via camera design
- Reduced by using a longer exposure time t
- Reduced by increasing the scene luminosity, higher P (e.g., using a flash)
- Reduced by increasing the aperture, higher P

Thermal noise

- Reduced by cooling the CCD, *i.e.*, lower D $\Rightarrow$ More expensive cameras
- Or by using a longer exposure time t

Digital optical imagery – How to reduce noise?

$$\text{SNR} = \frac{x}{\sqrt{x + \lambda + \sigma^2}} \quad \text{where} \quad x = PQ_e t, \quad \lambda = Dt$$

Photon noise

- Cannot be reduced via camera design
- Reduced by using a longer exposure time t
- Reduced by increasing the scene luminosity, higher P (e.g., using a flash)
- Reduced by increasing the aperture, higher P

Thermal noise

- Reduced by cooling the CCD, *i.e.*, lower D $\Rightarrow$ More expensive cameras
- Or by using a longer exposure time t

Readout noise

- Reduced by employing carefully designed electronics, *i.e.*, lower σ

  $\Rightarrow$ More expensive cameras

Or, reduced by image restoration softwares.

Digital optical imagery – Are these models accurate?

Processing pipeline

- There are always some pre-processing steps such as
  - white balance: to make sure neutral colors appear neutral,
  - demosaicing: to create a color image from incomplete color samples,
  - γ-correction: to optimize the usage of bits, and fit human perception of brightness,
  - compression: to improve memory usage (*e.g.*, JPEG).
- Technical details often hidden by the camera vendors.
- The noise in the resulting image becomes much harder to model.

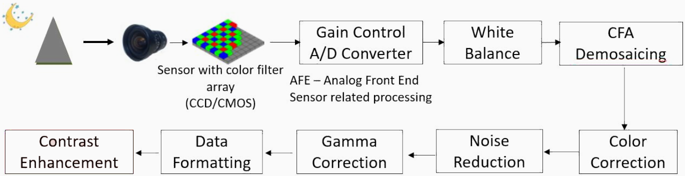

Source: Y. Gong and Y. Lee

Example (γ-correction)

$$y^{(\text{new})} = Ay^{\gamma}$$

(a) Non corrected (b) γ-corrected (c) Zoom $\times 8$ (d) Zoom $\times 30$

Gamma correction changes the nature of the noise. Since A and γ are usually not known, it becomes almost impracticable to model the noise accurately. In many scenarios, approximative models are used. The additive white Gaussian noise (AWGN) model is often considered for its simplicity.

Digital optical imagery – Noise models and post-processing

Example (Demosaicing)

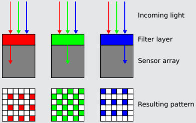

(a) Bayer filter

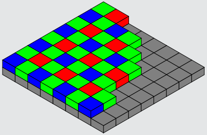

(b) Bayer pattern

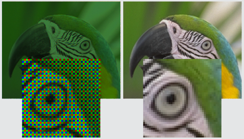

(c) Demosaicing

Basic idea:

- Use interpolation techniques.
- Bilinear interpolation: the red value of a non-red pixel is computed as the average of the two or four adjacent red pixels, and similarly for blue and green.

What is the influence on the noise?

- noise is no longer independent from one pixel to another,
- noise becomes spatially correlated.

Compression also creates spatial correlations.

Reminder of basic statistics

- X and Y two real random variables (*e.g.*, two pixel values)
- Independence: $p_{X,Y}(x, y) = p_X(x)p_Y(y)$
- Decorrelation: $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$
- Covariance:
$$\mathrm{Cov}(X, Y) = \mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])]$$
$$\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y] + \mathrm{Cov}(X, Y)$$
$$\mathrm{Var}(X) = \mathrm{Cov}(X, X)$$
- Correlation:
$$\mathrm{Corr}(X, Y) = \frac{\mathrm{Cov}(X, Y)}{\sqrt{\mathrm{Var}[X]\mathrm{Var}[Y]}}$$
$$\mathrm{Corr}(X, X) = 1$$

Reminder of basic statistics

- X and Y two real random variables (*e.g.*, two pixel values)
- Independence: $p_{X,Y}(x,y) = p_X(x)p_Y(y)$
- Decorrelation: $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$
- Covariance: $\mathrm{Cov}(X,Y) = \mathbb{E}[(X-\mathbb{E}[X])(Y-\mathbb{E}[Y])]$
  $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y] + \mathrm{Cov}(X,Y)$
  $\mathrm{Var}(X) = \mathrm{Cov}(X,X)$
- Correlation: $\mathrm{Corr}(X,Y) = \dfrac{\mathrm{Cov}(X,Y)}{\sqrt{\mathrm{Var}[X]\mathrm{Var}[Y]}}$
  $\mathrm{Corr}(X,X) = 1$

- Independence $\Leftrightarrow / \Rightarrow / \Leftarrow$ Decorrelation ?

Reminder of multivariate statistics

- $X = \begin{pmatrix} X_1 \\ X_2 \\ \vdots \\ X_n \end{pmatrix}$ and $Y = \begin{pmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_m \end{pmatrix}$ two real random vectors
- Entries are independent: $p_X(x) = \prod_k p_{X_k}(x_k)$
- Covariance matrix: $\mathrm{Var}(X) = \mathbb{E}[(X - \mathbb{E}[X])(X - \mathbb{E}[X])^T] \in \mathbb{R}^{n \times n}$

  $\mathrm{Var}(X)_{ij} = \mathrm{Cov}(X_i, X_j)$
- Correlation matrix $\mathrm{Corr}(X)_{ij} = \mathrm{Corr}(X_i, X_j)$
- Cross-covariance matrix: $\mathrm{Cov}(X, Y) = \mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])^T] \in \mathbb{R}^{n \times m}$
- Cross-correlation matrix: $\mathrm{Corr}(X, Y)_{ij} = \mathrm{Corr}(X_i, Y_j)$

Note: cross-correlation definition is slightly different in signal processing (in few slides)

Digital optical imagery – Noise models and correlations

- See an image x as a vector of $\mathbb{R}^n$,
- Its observation y is a realization of a random vector

$$Y = x + W.$$

- In general, noise is assumed to be zero-mean $\mathbb{E}[W] = 0$, then

$$\mathbb{E}[Y] = x \quad \text{and} \quad \mathrm{Var}[Y] = \mathrm{Var}[W] = \mathbb{E}[WW^T] = \boldsymbol{\Sigma}.$$

- $\boldsymbol{\Sigma}$ encodes variances and correlations (may depend on x).
- p_Y is often modeled with a multivariate Gaussian/normal distribution

$$p_Y(y; x) \approx \frac{1}{\sqrt{2\pi}^n |\boldsymbol{\Sigma}|^{1/2}} \exp\left(-\frac{1}{2}(y-x)^T \boldsymbol{\Sigma}^{-1}(y-x)\right).$$

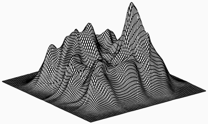

Underlying noise distribution

$\approx$

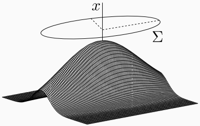

Gaussian approximation $Y \sim \mathcal{N}(x; \Sigma)$

Properties of covariance matrices

- $\mathbf{\Sigma} = \mathrm{Var}[Y]$ is square, symmetric and non-negative definite:

$$x^T \mathbf{\Sigma} x \geqslant 0, \quad \text{for all } x \neq 0 \text{ (eigenvalues } \lambda_i \geqslant 0).$$

- If all Y_k are linearly independent, then
  - $\mathbf{\Sigma}$ is positive definite: $x^T \mathbf{\Sigma} x > 0$, for all $x \neq 0$ ($\lambda_i > 0$),
  - $\mathbf{\Sigma}$ is invertible and $\mathbf{\Sigma}^{-1}$ is also symmetric positive definite,
  - Mahalanobis distance: $\sqrt{(y-x)^T \mathbf{\Sigma}^{-1} (y-x)} = \|\mathbf{\Sigma}^{-1/2}(y-x)\|_2$,
  - Its isoline $\left\{ y \; ; \; \|\mathbf{\Sigma}^{-1/2}(y-x)\|_2 = c, c > 0 \right\}$ describes an ellipsoid of center x and semi-axes the eigenvectors e_i with length $c\lambda_i$.

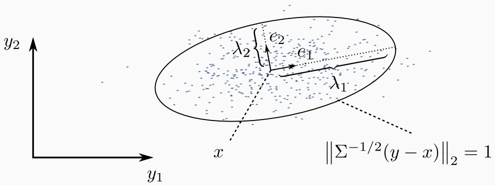

Digital optical imagery – Noise dictionary

Vocabulary in signal processing

- White noise: zero-mean noise + no correlations
- Stationary noise: identically distributed whatever the location
- Colored noise: stationary with pixels influencing their neighborhood
- Signal dependent: noise statistics depends on the signal intensity
- Space dependent: noise statistics depends on the location
- AWGN: Additive White Gaussian Noise: $Y \sim \mathcal{N}(x; \sigma^2 \mathrm{Id}_n)$

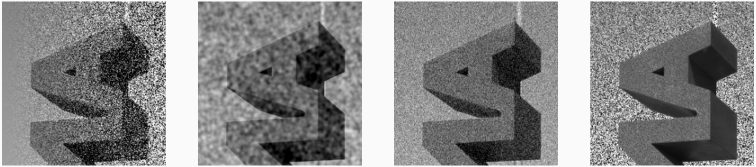

How is it encoded in Σ?

1. Σ diagonal: noise is uncorrelated – *white*
2. $\Sigma_{i,i} = f(s_i)$: variance depends on pixel location s_i – *space dependent*
3. $\Sigma_{i,i} = f(x_i)$: variance depends on pixel value x_i – *signal dependent*
4. $\Sigma_{i,j} = f(s_i - s_j)$: correlations depends on the shift – *stationary*

   For 1d signals, Σ is Toeplitz: $\Sigma = \begin{pmatrix} a & b & \dots & c \\ d & a & \ddots & \vdots \\ \vdots & \ddots & \ddots & b \\ e & \dots & d & a \end{pmatrix}$

5. $\Sigma = \underbrace{\begin{pmatrix} \sigma^2 & & 0 \\ & \ddots & \\ 0 & & \sigma^2 \end{pmatrix}}_{=\sigma^2 \mathrm{Id}_n}$: noise is homoscedastic ($\neq$ heteroscedastic) – *white+stationary*

Digital optical imagery – Settings to avoid noise

(a) Very short exposure

(b) Short exposure

(c) Flash

(d) Normal exposure

(e) Long exposure

(f) Long + hand shaking

- Short exposure: too much noise
- Using a flash: change the aspect of the scene
- Long exposure: subject to blur and saturation (use a tripod)

What is blur?

Blur: The best of, 2000

Digital optical imagery – Blur

Motion blur

- Moving object
- Camera shake
- Atmospheric turbulence
- Long exposure time

Camera blur

- Limited resolution
- Diffraction
- Bad focus
- Wrong optical design

Bokeh

- Out-of-focus parts
- Often for artistic purpose

Motion blur

London (UK)

Munich (Germany)

Camera blur

Hubble Space Telescope (NASA)

Bokeh

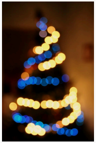

Christmas tree

Mulholand drive (2001)

Digital optical imagery – Blur

Motion blur

- Moving object
- Camera shake
- Atmospheric turbulence
- Long exposure time

Camera blur

- Limited resolution
- Diffraction
- Bad focus
- Wrong optical design

Bokeh

- Out-of-focus parts
- Often for artistic purpose

Motion blur

London (UK)

Munich (Germany)

Camera blur

Hubble Space Telescope (NASA)

Bokeh

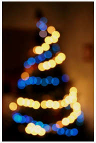

Christmas tree

Mulholand drive (2001)

How to model blur?

Linear model of blur

- Observed pixel values are a mixture of the underlying ones

$$y_{i,j} = \sum_{k=1}^{n} \sum_{l=1}^{n} h_{i,j,k,l} x_{k,l} \quad \text{where} \quad h_{k,l} \geqslant 0 \text{ and } \sum_{l=1}^{n} h_{k,l} = 1$$

- Matrix/vector representation: $y = \boldsymbol{H} x \qquad y \in \mathbb{R}^n$, $x \in \mathbb{R}^n$, $\boldsymbol{H} \in \mathbb{R}^{n \times n}$

Digital optical imagery – Linear blur

Linear model of blur

- Observed pixel values are a mixture of the underlying ones

$$y_{i,j} = \sum_{k=1}^{n} \sum_{l=1}^{n} h_{i,j,k,l} x_{k,l} \quad \text{where} \quad h_{k,l} \geqslant 0 \text{ and } \sum_{l=1}^{n} h_{k,l} = 1$$

- Matrix/vector representation: $y = \boldsymbol{H} x \qquad y \in \mathbb{R}^n,\ x \in \mathbb{R}^n,\ \boldsymbol{H} \in \mathbb{R}^{n \times n}$

$$y = \left(\begin{array}{ccc|c|ccc}
\multicolumn{3}{c}{\overbrace{\qquad\qquad\qquad}^{\text{First line}}} & & \multicolumn{3}{c}{\overbrace{\qquad\qquad\qquad}^{\text{Last line}}} \\
h_{1,1,1,1} & \dots & h_{1,1,1,n_2} & \dots & h_{1,1,n_1,1} & \dots & h_{1,1,n_1,n_2} \\
\vdots & & \vdots & & \vdots & & \vdots \\
h_{1,n_2,1,1} & \dots & h_{1,n_2,1,n_2} & \dots & h_{1,n_2,n_1,1} & \dots & h_{1,n_2,n_1,n_2} \\
\hline
\vdots & & \vdots & & \vdots & & \vdots \\
\hline
h_{n_1,1,1,1} & \dots & h_{n_1,1,1,n_2} & \dots & h_{n_1,1,n_1,1} & \dots & h_{n_1,1,n_1,n_2} \\
\vdots & & \vdots & & \vdots & & \vdots \\
h_{n_1,n_2,1,1} & \dots & h_{n_1,n_2,1,n_2} & \dots & h_{n_1,n_2,n_1,1} & \dots & h_{n_1,n_2,n_1,n_2}
\end{array}\right)
\left(\begin{array}{c}
x_{1,1} \\ \vdots \\ x_{1,n_2} \\ \hline \vdots \\ \hline x_{n_1,1} \\ \vdots \\ x_{n_1,n_2}
\end{array}\right)$$

Digital optical imagery – Point Spread Function (PSF)

$$
\boldsymbol{H} \times \underbrace{\begin{pmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}}_{\substack{\text{Only one 1 for} \\ \text{some pixel } (i,j)}} = \underbrace{\begin{pmatrix} h_{1,1,i,j} \\ \vdots \\ h_{i,j-1,i,j} \\ h_{i,j,i,j} \\ h_{i,j+1,i,j} \\ \vdots \\ \vdots \\ h_{n_1,n_2,i,j} \end{pmatrix}}_{\text{One column of } \boldsymbol{H}} \xrightarrow{\text{“reshape”}} \underbrace{\begin{pmatrix} h_{1,1,i,j} & h_{2,1,i,j} & \dots & h_{n_1,1,i,j} \\ \vdots & \vdots & & \vdots \\ h_{1,n_2,i,j} & h_{2,n_2,i,j} & \dots & h_{n_1,n_2,i,j} \end{pmatrix}}_{\substack{\text{System's impulse response at location } (i,j) \\ \text{called, Point spread function}}}
$$

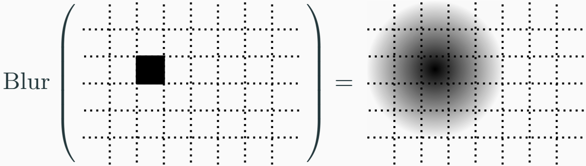

Spatially varying PSF – non-stationary blur

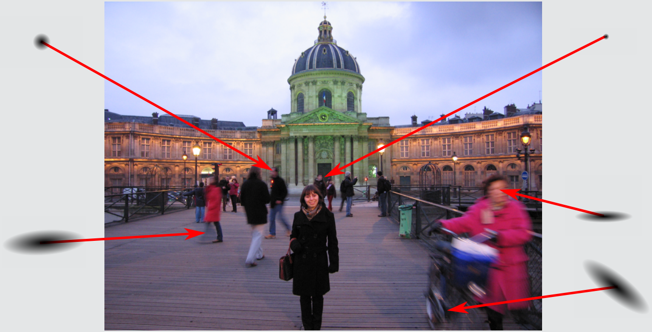

Digital optical imagery – Stationary blur

Stationary blur

- Shift invariant: blurring depends only on the relative position:

$$h_{i,j,k,l} = \kappa_{k-i,l-j},$$

  i.e., same PSF everywhere.

- Corresponds to the (discrete) cross-correlation *(not the same as in statistics)*

$$y = \kappa \star x \quad \Leftrightarrow \quad y_{i,j} = \sum_{k=-\infty}^{+\infty} \sum_{l=-\infty}^{+\infty} \kappa_{k,l} x_{i+k,j+l}$$

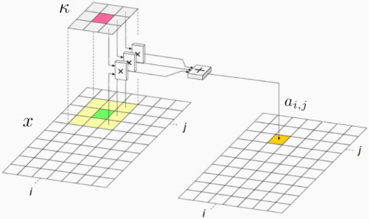

Stationary blur

- Shift invariant: blurring depends only on the relative position:

$$h_{i,j,k,l} = \kappa_{k-i,l-j},$$

i.e., same PSF everywhere.

- Corresponds to the (discrete) cross-correlation *(not the same as in statistics)*

$$y = \kappa \star x \quad \Leftrightarrow \quad y_{i,j} = \sum_{k=-\infty}^{+\infty} \sum_{l=-\infty}^{+\infty} \kappa_{k,l} x_{i+k,j+l}$$

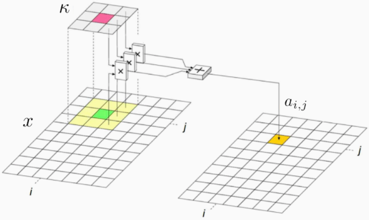

Here κ has a $q = 3 \times 3$ support

$$\Rightarrow \sum_{k=-\infty}^{+\infty} \sum_{l=-\infty}^{+\infty} \equiv \sum_{k=-1}^{+1} \sum_{l=-1}^{+1}$$

q called window size.

Direct computation requires $O(nq)$.

$$\Rightarrow q \ll n$$

Cross-correlation vs Convolution product

- If κ is complex then the cross-correlation becomes

$$y = \kappa \star x \quad \Leftrightarrow \quad y_{i,j} = \sum_{k=-\infty}^{+\infty} \sum_{l=-\infty}^{+\infty} \kappa_{k,l}^* x_{i+k,j+l}.$$

- Complex conjugate: $(a + ib)^* = a - ib$.
- $y = \kappa \star x$ can be re-written as the (discrete) convolution product

$$y = \nu * x \quad \Leftrightarrow \quad y_{i,j} = \sum_{k=-\infty}^{+\infty} \sum_{l=-\infty}^{+\infty} \nu_{k,l} x_{i-k,j-l} \quad \text{with} \quad \nu_{k,l} = \kappa_{-k,-l}^*.$$

- ν called convolution kernel.

Why convolution instead of cross-correlation?

- **Associative:** $(f * g) * h = f * (g * h)$
- **Commutative:** $f * g = g * f$

For cross-correlation, only true if the signal is Hermitian, i.e., if $f_{k,l} = f_{-k,-l}^*$.

3×3 box convolution

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

	0								

Source: Steven Seitz

3×3 box convolution

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

	0	10							

Source: Steven Seitz

3×3 box convolution

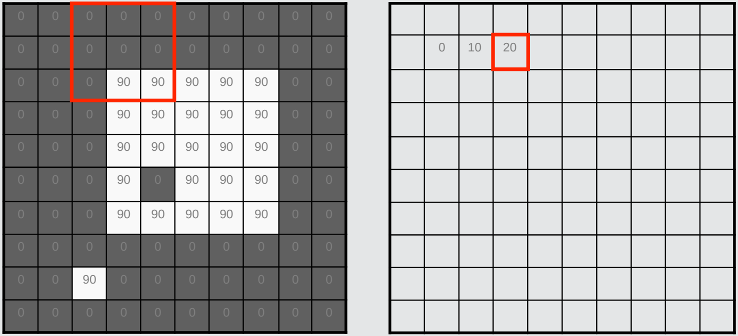

Source: Steven Seitz

3×3 box convolution

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

	0	10	20	30					

Source: Steven Seitz

3×3 box convolution

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

	0	10	20	30	30				

Source: Steven Seitz

3×3 box convolution

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

	0	10	20	30	30	30	20	10	
	0	20	40	60	60	60	40	20	
	0	30	60	90	90	90	60	30	
	0	30	50	80	80	90	60	30	
	0	30	50	80	80	90	60	30	
	0	20	30	50	50	60	40	20	
	10	20	30	30	30	30	20	10	
	10	10	10	0	0	0	0	0	

Source: Steven Seitz

Classical kernels

- Box kernel:

$$\kappa_{i,j} = \frac{1}{Z} \begin{cases} 1 & \text{if} \quad \max(|i|,|j|) \leqslant \tau \\ 0 & \text{otherwise} \end{cases}$$

- Gaussian kernel:

$$\kappa_{i,j} = \frac{1}{Z} \exp\left(-\frac{i^2 + j^2}{2\tau^2} \right)$$

- Exponential kernel:

$$\kappa_{i,j} = \frac{1}{Z} \exp\left(-\frac{\sqrt{i^2 + j^2}}{\tau} \right)$$

- Z normalization constant s.t. $\sum_{i,j} \kappa_{i,j} = 1$

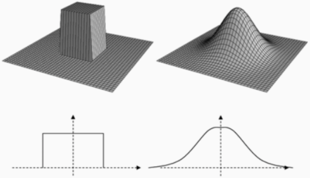

0	0	0	0	0
0	1	1	1	0
0	1	1	1	0
0	1	1	1	0
0	0	0	0	0

0	1	2	1	0
1	3	5	3	1
2	5	9	5	2
1	3	5	3	1
0	1	2	1	0

Digital optical imagery – Gaussian kernel

$$\kappa_{i,j} = \frac{1}{Z} \exp\left(-\frac{i^2 + j^2}{2\tau^2}\right)$$

Influence of τ

- $\sqrt{i^2 + j^2}$: distance to the central pixel,
- τ: controls the influence of neighbor pixels, *i.e.*, the strength of the blur

Small τ

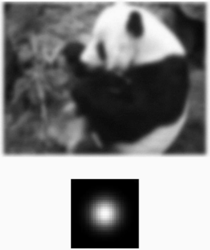

Medium τ

...

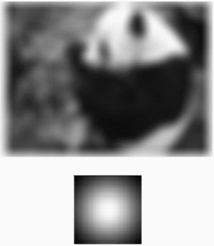

Large τ

How to deal when the kernel window overlaps outside the image domain?

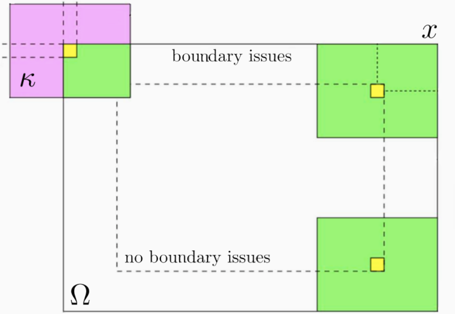

i.e., how to evaluate $y_{i,j} = \sum_{k,l} \kappa_{k,l} x_{i+k,j+l}$ when $(i+k, j+l) \notin \Omega$?

Standard techniques:

zero-padding

extension

mirror

periodical

Other common problems

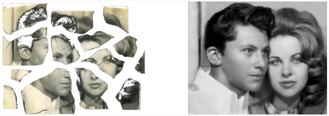

Source: Wikipedia

Digital optical imagery – Other "standard" noise models

Transmission, encoding, compression, rendering can lead to other models of corruptions assimilated to noise.

Salt-and-pepper noise

- Randomly saturated pixels to black (value 0) or white (value $L-1$)

$$p_Y(y\,;\,x) = \begin{cases} 1-P & \text{if } y = x \\ P/2 & \text{if } y = 0 \\ P/2 & \text{if } y = L-1 \\ 0 & \text{otherwise} \end{cases}$$

$P = 10\%$

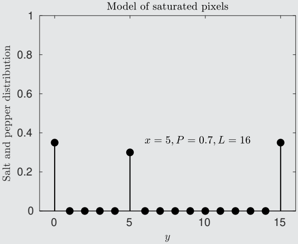

Impulse noise

- Some pixels take "arbitrary" values

$$p_Y(y\,;\,x) = \begin{cases} 1 - P + P/L & \text{if } y = x \\ P/L & \text{otherwise} \end{cases}$$

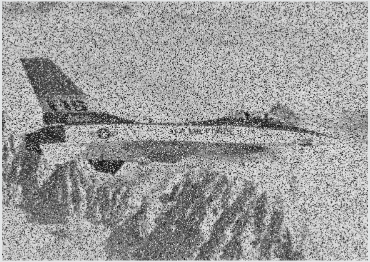

$P = 40\%$

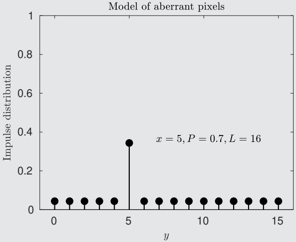

(other models exist: Laplacian, Cauchy, . . .)

Digital optical imagery – Corruptions assimilated to noise

Corruptions assimilated to noise

- compression artifacts,
- data corruption,
- rendering (*e.g.*, half-toning).

(a) Source image

$\longrightarrow$

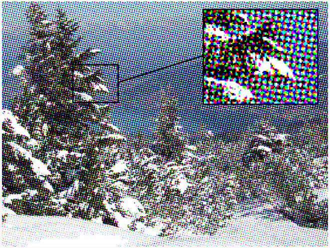

(b) Half-toned image

Deconvolution subject to noise

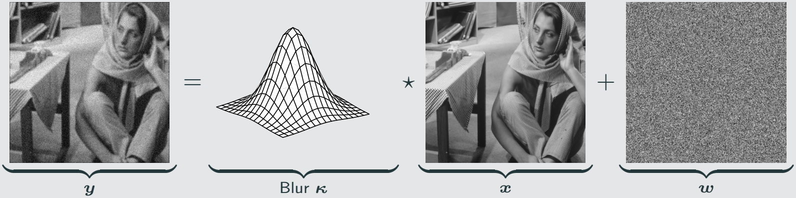

Goal: Retrieve the sharp and clean image x from y

Inpainting (mask)

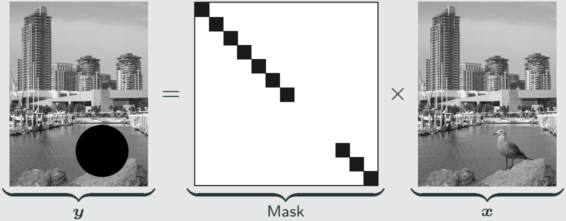

Goal: Fill the hole

Single-frame super-resolution (sub-sampling + convolution + noise)

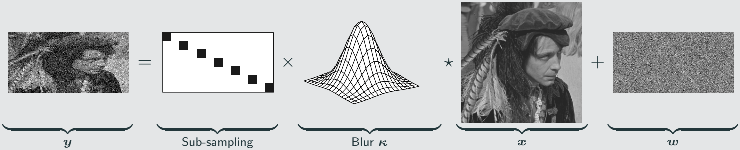

Goal: Increase the resolution of the Low Resolution (LR) image $\boldsymbol{y}$ to retrieve the High Resolution (HR) image $\boldsymbol{x}$

Multi-frame super-resolution (different sub-pixel shifts + noise)

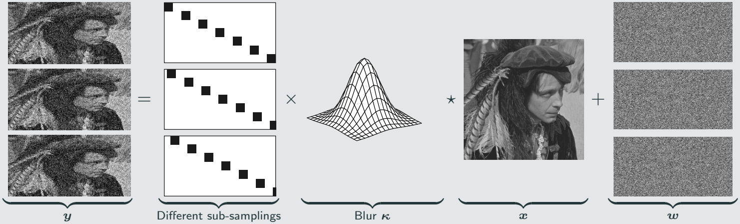

Goal: Combine the information of LR images $\boldsymbol{y}_k$ to retrieve the HR image $\boldsymbol{x}$

Compressed sensing

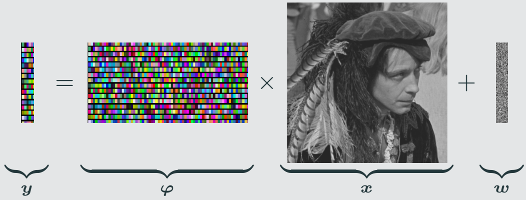

- Goal: compress the quantity of information,
  e.g., to reduce acquisition time or transmission cost,
  and provide guarantee to reconstruct or approximate $\boldsymbol{x}$.

- Unlike classical compression techniques (jpeg, . . .):
  - no compression steps,
  - sensor designed to provide directly the coefficients $\boldsymbol{y}$,
  - the decompression time is usually not an issue.

Digital optical imagery – Other sources of corruptions

- Quantization
- Saturation
- Aliasing
- Compression artifacts
- Chromatic aberrations
- Dead/Stuck/Hot pixels

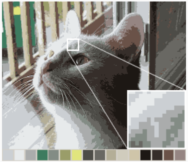

(a) 4-bit quantization

vs

(b) Saturation (overexposure)

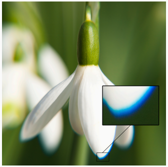

(c) Color aberrations

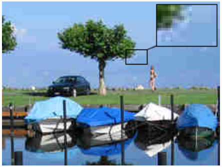

(d) Compression artifacts

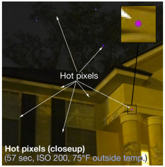

(e) Hot pixels

Sources: Wikipedia, David C. Pearson, Dpreview

Digital optical imagery – A technique to avoid saturation

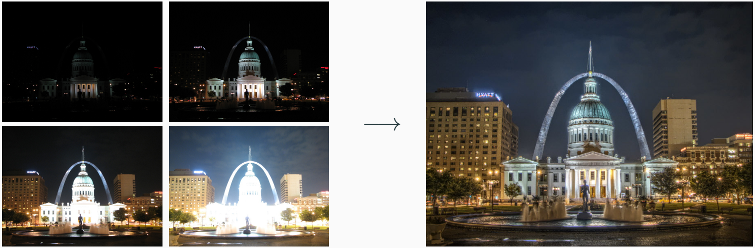

Figure 3 – Fusion of under- and over-exposed images (St Louis, Missouri, USA)

High dynamic range imaging

- Goal: avoid saturation effects
- Technique: merge several images with different exposure times
- Tone mapping: problem of displaying an HDR image on a screen
- Remark: there also exist HDR sensors

Digital optical imagery – Why chromatic aberrations?

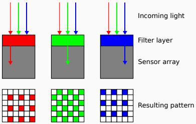

(a) Bayer filter

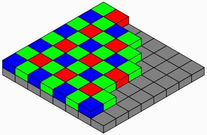

(b) Bayer pattern

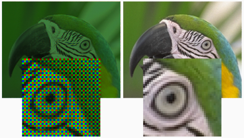

(c) Demosaicing

(d) Results of different algorithms (Source: DMMD)

Demosaicing

- Goal: reconstruct a color image from the incomplete color samples
- Problem: standard interpolation techniques lead to chromatic aberrations

Non-conventional imagery

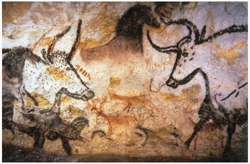

Depiction of aurochs, horses and deer (Lascaux, France)

Passive versus active imagery

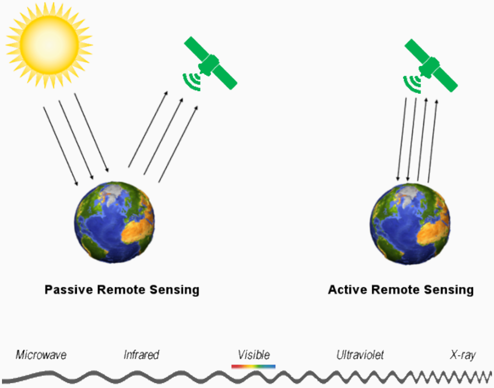

- Passive: optical (visible), infrared, hyper-spectral (several frequencies).
- Active: radar (microwave), sonar (radio), CT scans (X-ray), MRI (radio).

Synthetic aperture radar (SAR) imagery

Synthetic aperture radar (SAR) imaging systems

- Mounted on an aircraft or spacecraft,
- Measures echoes of a back-scattered electromagnetic wave (microwave),
- Signal carries information about geophysical properties of the scene,
- Used for earth monitoring and military surveillance,
  - deforestation, flooding, urban growth, earthquake, glaciology, . . .
- Performs day and night and in any weather conditions.

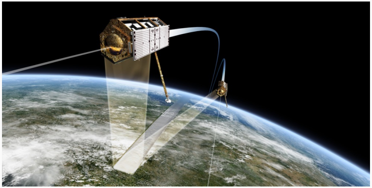

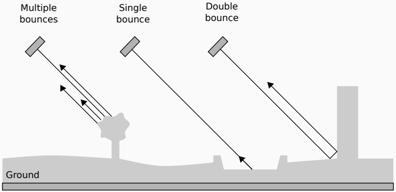

Synthetic aperture radar (SAR) imagery

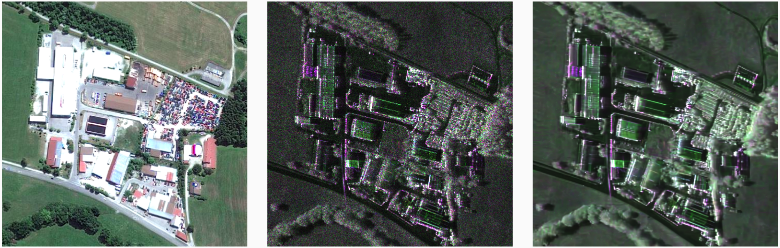

(a) Optical (b) SAR (c) Denoising result

SAR images are corrupted by speckle

- Source of fluctuation: arbitrary roughness/rugosity of the scene
- Magnitude $y \in \mathbb{R}^+$ fluctuates around its means $x \in \mathbb{R}^+$
- Fluctuations proportional to x
- Gamma distributed
- Multiplicative behavior: $y = x \times s$
- Signal dependent with constant SNR

Other examples of speckle

Sonar imagery

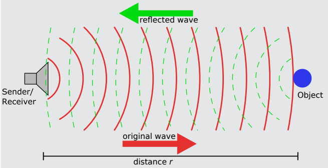

Submerged plane wreckage

Ultrasound imagery

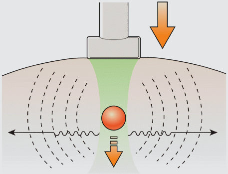

Ultrasound image of a fetus

Computed tomography (CT) imaging systems

- Uses irradiations to scan a 3d volume
- Measures attenuations in several directions
- Runs a 3d reconstruction algorithm

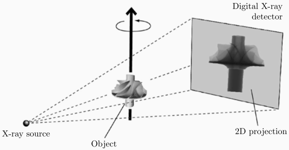

- Industry
  - Defect analysis
  - Computer-aided design
  - Material analysis
  - Petrophysics
  - . . .

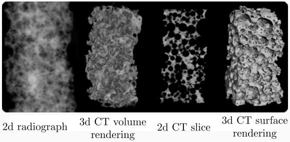

- Medical imagery
  - X-ray CT
  - Positron emission tomography (PET)
  - Medical diagnoses
  - . . .

Shot noise

- Due to the limited number of X-ray photons reaching the detector,
- Poisson distributed,
- SNR increases with exposure time,
- Higher exposure $\Rightarrow$ higher irradiation ☹.

Computed tomography (CT) imaging systems

Shot noise

- Due to the limited number of X-ray photons reaching the detector,
- Poisson distributed,
- SNR increases with exposure time,
- Higher exposure ⇒ higher irradiation ☹.

Streaking

- Due to the limited number of projection angles,
- Linear degradation model: $y = \boldsymbol{H}x$,
- More projections ⇒ better reconstruction ☺, but higher irradiation ☹.

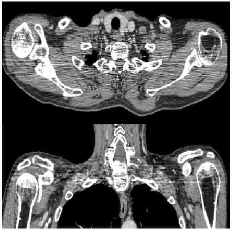

Magnetic resonance imaging (MRI)

- Apply a strong magnetic field varying along the patient (gradient),
- Hydrogen nucleus' spins align with the field,
- Emit a pulse to change the alignments of spins in a given slice,
- Nuclei return to equilibrium: measure its released radio frequency signal,
- Repeat for the different slices by applying different frequency pulses,
- Use algorithms to reconstruct a 3d volume from raw signals.

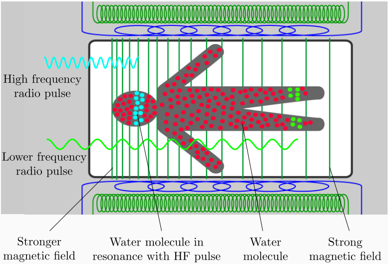

Magnetic resonance imaging (MRI)

- Apply a strong magnetic field varying along the patient (gradient),
- Hydrogen nucleus' spins align with the field,
- Emit a pulse to change the alignments of spins in a given slice,
- Nuclei return to equilibrium: measure its released radio frequency signal,
- Repeat for the different slices by applying different frequency pulses,
- Use algorithms to reconstruct a 3d volume from raw signals.

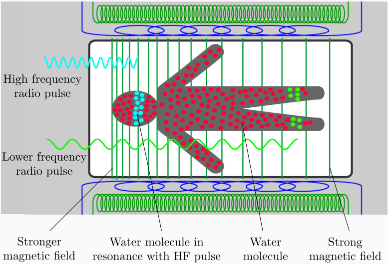

Unlike CT scans, no harmful radiation!

Rician noise

- Main source of noise: thermal motions in patient's body emit radio waves
- Magnitude $y \in \mathbb{R}^+$ fluctuates (for x large enough) around: $\sqrt{x^2 + \sigma^2}$
- Fluctuations approximately equal (for x large enough) to σ^2
- Rician distributed

Streaking: due to limited number of acquisitions

- As in CT scans, linear corruptions: $y = \boldsymbol{H}x$.

$\Rightarrow$ using a longer acquisition time, but limited by
- cost,
- patient comfort.

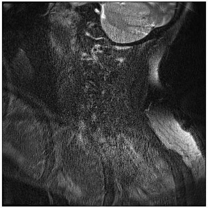

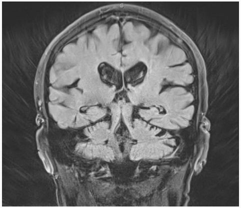

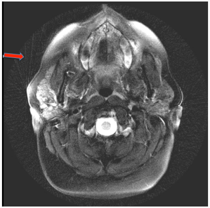

Major image restoration issues

Jacques Hadamard (1865–1963)

Major image restoration issues

Usual image degradation models

- Images often viewed through a linear operator (*e.g.*, blur or streaking)

$$y = \boldsymbol{H}x \quad \Leftrightarrow \quad \begin{cases} h_{11}x_1 + h_{12}x_2 + \ldots + h_{1n}x_n & = y_1 \\ h_{21}x_1 + h_{22}x_2 + \ldots + h_{2n}x_n & = y_2 \\ \vdots & \\ h_{n1}x_1 + h_{n2}x_2 + \ldots + h_{nn}x_n & = y_n \end{cases}$$

- Retrieving $x \Rightarrow$ Inverting $\boldsymbol{H}$ (*i.e.*, solving the system of linear equations)

$$\hat{x} = \boldsymbol{H}^{-1}y$$

(a) Unknown image x (b) Observation y (c) Estimate $\hat{x}$

Is image restoration solved then?

Major image restoration issues

Limitations

- $\boldsymbol{H}$ is often non-invertible
  - equations are linearly dependent,
  - system is under-determined,
  - infinite number of solutions,
  - which one to choose?
- The system is said to be ill-posed in opposition to well-posed.

Well-posed problem (Hadamard)

1. a solution exists,
2. the solution is unique,
3. the solution's behavior changes continuously with the initial conditions.

Major image restoration issues

Limitations

- Or, $\boldsymbol{H}$ is invertible but ill-conditioned:
  - small perturbations in y lead to large errors in $\hat{x} = \boldsymbol{H}^{-1}y$,
  - and unfortunately y is often corrupted by noise: $y = \boldsymbol{H}x + w$,
  - and unfortunately y is often encoded with limited precision.

(a) Unknown image x (b) Observation y (c) Estimate $\hat{x}$

- Condition-number: $\kappa(\boldsymbol{H}) = \|\boldsymbol{H}^{-1}\|_2 \|\boldsymbol{H}\|_2 = \frac{\sigma_{\max}}{\sigma_{\min}}$

(σ_k singular values of $\boldsymbol{H}$)

- the larger $\kappa(\boldsymbol{H}) \geqslant 1$, the more ill-conditioned/difficult is the inversion.

Questions?

Next class: histogram manipulation and basics of filtering

Slides from Charles Deledalle

Sources, images courtesy and acknowledgment

L. Condat
DLR
DMMD
Dpreview
Y. Gong
A. Horodniceanu

I. Kokkinos
J.-M. Nicolas
A. Newson
D. C. Pearson
S. Seitz

V. Tong Ta
P. Tilakaratna
Wikipedia
R. Willett
Y. Lee