Digital Image Processing

Ming Jiang

School of Mathematical Sciences
Peking University

ming-jiang@pku.edu.cn

May 18, 2010

Outline I

Basic concepts
- Image functions
- Image formation
- Image as a stochastic process

Image digitization
- Sampling
- Quantization

Digital image properties
- Metric and topological properties of digital images
  - Metric properties of digital images
  - Topological properties of digital images
  - Contiguity paradoxes
  - Other topological and geometrical properties
- Histogram
- Entropy
- Visual perception of the image
  - Contrast
  - Acuity
  - Visual Illusions

Outline II

- Image quality
- Noise in images
  - Image Noise
  - Noise Type
  - Simulation of noise
  - The rejection method
  - The Polar method
  - Poissonian noise
  - Mixed noise: Poissonian + Gaussian

- Color Images
  - Physics of Color
  - Color Perceived by Humans
  - Color Spaces
  - Palette Images
  - Color Constancy

Signals

- Fundamental concepts and mathematical tools are introduced in this chapter which will be used throughout the course.
- Mathematical models are often used to describe images and other signals.
- A signal is a function depending on some variable(s) with physical meaning. Signals can be
  - one-dimensional (e.g., audio signal dependent on time);
  - two-dimensional (e.g., images dependent on two co-ordinates in a plane);
  - three-dimensional (e.g., describing an object in space or video signal);
  - or higher-dimensional.

Signals

- Fundamental concepts and mathematical tools are introduced in this chapter which will be used throughout the course.
- Mathematical models are often used to describe images and other signals.
- A signal is a function depending on some variable(s) with physical meaning. Signals can be
  - one-dimensional (e.g., audio signal dependent on time);
  - two-dimensional (e.g., images dependent on two co-ordinates in a plane);
  - three-dimensional (e.g., describing an object in space or video signal);
  - or higher-dimensional.

Signals

- Fundamental concepts and mathematical tools are introduced in this chapter which will be used throughout the course.
- Mathematical models are often used to describe images and other signals.
- A signal is a function depending on some variable(s) with physical meaning. Signals can be
  - one-dimensional (e.g., audio signal dependent on time);
  - two-dimensional (e.g., images dependent on two co-ordinates in a plane);
  - three-dimensional (e.g., describing an object in space or video signal);
  - or higher-dimensional.

Signals

- Fundamental concepts and mathematical tools are introduced in this chapter which will be used throughout the course.
- Mathematical models are often used to describe images and other signals.
- A signal is a function depending on some variable(s) with physical meaning. Signals can be
  - one-dimensional (e.g., audio signal dependent on time);
  - two-dimensional (e.g., images dependent on two co-ordinates in a plane);
  - three-dimensional (e.g., describing an object in space or video signal);
  - or higher-dimensional.

Signals

- Fundamental concepts and mathematical tools are introduced in this chapter which will be used throughout the course.
- Mathematical models are often used to describe images and other signals.
- A signal is a function depending on some variable(s) with physical meaning. Signals can be
  - one-dimensional (e.g., audio signal dependent on time);
  - two-dimensional (e.g., images dependent on two co-ordinates in a plane);
  - three-dimensional (e.g., describing an object in space or video signal);
  - or higher-dimensional.

Signals

- Fundamental concepts and mathematical tools are introduced in this chapter which will be used throughout the course.
- Mathematical models are often used to describe images and other signals.
- A signal is a function depending on some variable(s) with physical meaning. Signals can be
  - one-dimensional (e.g., audio signal dependent on time);
  - two-dimensional (e.g., images dependent on two co-ordinates in a plane);
  - three-dimensional (e.g., describing an object in space or video signal);
  - or higher-dimensional.

Signals

- Fundamental concepts and mathematical tools are introduced in this chapter which will be used throughout the course.
- Mathematical models are often used to describe images and other signals.
- A signal is a function depending on some variable(s) with physical meaning. Signals can be
  - one-dimensional (e.g., audio signal dependent on time);
  - two-dimensional (e.g., images dependent on two co-ordinates in a plane);
  - three-dimensional (e.g., describing an object in space or video signal);
  - or higher-dimensional.

Images as 2D functions: discussions

- A scalar function may be sufficient to describe a monochromatic image.
- Vector functions are to represent, for example, color images consisting of three component colors.

- Functions we shall work with may be categorized as continuous, discrete and digital.
- A continuous function has continuous domain and range;
- If the domain set is discrete, then we get a `discrete function`;
- if the range set is also discrete, then we have a `digital function`.

Images as 2D functions: discussions

- A scalar function may be sufficient to describe a monochromatic image.
- Vector functions are to represent, for example, color images consisting of three component colors.

- Functions we shall work with may be categorized as continuous, discrete and digital.
- A continuous function has continuous domain and range;
- If the domain set is discrete, then we get a `discrete function`;
- if the range set is also discrete, then we have a `digital function`.

Images as 2D functions: discussions

- A scalar function may be sufficient to describe a monochromatic image.
- Vector functions are to represent, for example, color images consisting of three component colors.

- Functions we shall work with may be categorized as continuous, discrete and digital.
- A continuous function has continuous domain and range;
- If the domain set is discrete, then we get a `discrete function`;
- if the range set is also discrete, then we have a `digital function`.

Images as 2D functions: discussions

- A scalar function may be sufficient to describe a monochromatic image.
- Vector functions are to represent, for example, color images consisting of three component colors.

- Functions we shall work with may be categorized as continuous, discrete and digital.
- A continuous function has continuous domain and range;
- If the domain set is discrete, then we get a `discrete function`;
- if the range set is also discrete, then we have a `digital function`.

Images as 2D functions: discussions

- A scalar function may be sufficient to describe a monochromatic image.
- Vector functions are to represent, for example, color images consisting of three component colors.

- Functions we shall work with may be categorized as continuous, discrete and digital.
- A continuous function has continuous domain and range;
- If the domain set is discrete, then we get a `discrete function`;
- if the range set is also discrete, then we have a `digital function`.

Images as 2D functions: discussions

- A scalar function may be sufficient to describe a monochromatic image.
- Vector functions are to represent, for example, color images consisting of three component colors.

- Functions we shall work with may be categorized as continuous, discrete and digital.
- A continuous function has continuous domain and range;
- If the domain set is discrete, then we get a `discrete function`;
- if the range set is also discrete, then we have a `digital function`.

Outline

Images as 2D functions: discussions

- An image can be modeled by a continuous function of two or three variables;
- Its arguments are co-ordinates x and y in a plane;
- If images change in time a third variable t might be added.
- The image function values correspond to the brightness at image points.

Images as 2D functions: discussions

- An image can be modeled by a continuous function of two or three variables;
- Its arguments are co-ordinates x and y in a plane;
- If images change in time a third variable t might be added.
- The image function values correspond to the brightness at image points.

Images as 2D functions: discussions

- An image can be modeled by a continuous function of two or three variables;
- Its arguments are co-ordinates x and y in a plane;
- If images change in time a third variable t might be added.
- The image function values correspond to the brightness at image points.

Images as 2D functions: discussions

- An image can be modeled by a continuous function of two or three variables;
- Its arguments are co-ordinates *x* and *y* in a plane;
- If images change in time a third variable *t* might be added.
- The image function values correspond to the brightness at image points.

Intepretation of image values

- The function value can express other physical quantities as well (temperature, pressure distribution, distance from the observer, etc.).
- The brightness integrates different optical quantities — using brightness as a basic quantity allows us to avoid the description of the very complicated process of image formation.

Intepretation of image values

- The function value can express other physical quantities as well (temperature, pressure distribution, distance from the observer, etc.).
- The brightness integrates different optical quantities — using brightness as a basic quantity allows us to avoid the description of the very complicated process of image formation.

Homework

Discuss the various factors that influence the brightness of a pixel in an image.

2D images

- The image on the human eye retina or on a TV camera sensor is intrinsically 2D.
- We shall call such a 2D image bearing information about brightness points an `intensity image`.
- The real world which surrounds us is intrinsically 3D.
- The 2D intensity image is the result of a perspective projection of the 3D scene.

2D images

- The image on the human eye retina or on a TV camera sensor is intrinsically 2D.
- We shall call such a 2D image bearing information about brightness points an `intensity image`.
- The real world which surrounds us is intrinsically 3D.
- The 2D intensity image is the result of a perspective projection of the 3D scene.

2D images

- The image on the human eye retina or on a TV camera sensor is intrinsically 2D.
- We shall call such a 2D image bearing information about brightness points an `intensity image`.
- The real world which surrounds us is intrinsically 3D.
- The 2D intensity image is the result of a perspective projection of the 3D scene.

2D images

- The image on the human eye retina or on a TV camera sensor is intrinsically 2D.
- We shall call such a 2D image bearing information about brightness points an `intensity image`.
- The real world which surrounds us is intrinsically 3D.
- The 2D intensity image is the result of a perspective projection of the 3D scene.

Ill-posed problem I

- When 3D objects are mapped into the camera plane by perspective projection a lot of information disappears as such a transformation is not one-to-one.
- Recognizing or reconstructing objects in a 3D scene from one image is an ill-posed problem.
- Recovering information lost by perspective projection is only one, mainly geometric, problem of computer vision.
- The aim is to recover a full 3D representation such as may be used in computer graphics.

Ill-posed problem I

- When 3D objects are mapped into the camera plane by perspective projection a lot of information disappears as such a transformation is not one-to-one.
- Recognizing or reconstructing objects in a 3D scene from one image is an ill-posed problem.
- Recovering information lost by perspective projection is only one, mainly geometric, problem of computer vision.
- The aim is to recover a full 3D representation such as may be used in computer graphics.

Ill-posed problem I

- When 3D objects are mapped into the camera plane by perspective projection a lot of information disappears as such a transformation is not one-to-one.
- Recognizing or reconstructing objects in a 3D scene from one image is an ill-posed problem.
- Recovering information lost by perspective projection is only one, mainly geometric, problem of computer vision.
- The aim is to recover a full 3D representation such as may be used in computer graphics.

Ill-posed problem I

- When 3D objects are mapped into the camera plane by perspective projection a lot of information disappears as such a transformation is not one-to-one.
- Recognizing or reconstructing objects in a 3D scene from one image is an ill-posed problem.
- Recovering information lost by perspective projection is only one, mainly geometric, problem of computer vision.
- The aim is to recover a full 3D representation such as may be used in computer graphics.

Ill-posed problem II

- The second problem is how to understand image brightness.
- The only information available in an intensity image is the brightness of pixels.
- They are dependent on a number of independent factors such as
  - object surface reflectance properties (given by the surface material, micro-structure and marking),
  - illumination properties,
  - object surface orientation with respect to a viewer and light source.
- This is a non-trivial and again ill-posed problem.

Ill-posed problem II

- The second problem is how to understand image brightness.
- The only information available in an intensity image is the brightness of pixels.
- They are dependent on a number of independent factors such as
  - object surface reflectance properties (given by the surface material, micro-structure and marking),
  - illumination properties,
  - object surface orientation with respect to a viewer and light source.
- This is a non-trivial and again ill-posed problem.

Ill-posed problem II

- The second problem is how to understand image brightness.
- The only information available in an intensity image is the brightness of pixels.
- They are dependent on a number of independent factors such as
  - object surface reflectance properties (given by the surface material, micro-structure and marking),
  - illumination properties,
  - object surface orientation with respect to a viewer and light source.
- This is a non-trivial and again ill-posed problem.

Ill-posed problem II

- The second problem is how to understand image brightness.
- The only information available in an intensity image is the brightness of pixels.
- They are dependent on a number of independent factors such as
  - object surface reflectance properties (given by the surface material, micro-structure and marking),
  - illumination properties,
  - object surface orientation with respect to a viewer and light source.
- This is a non-trivial and again ill-posed problem.

Ill-posed problem II

- The second problem is how to understand image brightness.
- The only information available in an intensity image is the brightness of pixels.
- They are dependent on a number of independent factors such as
  - object surface reflectance properties (given by the surface material, micro-structure and marking),
  - illumination properties,
  - object surface orientation with respect to a viewer and light source.
- This is a non-trivial and again ill-posed problem.

Ill-posed problem II

- The second problem is how to understand image brightness.
- The only information available in an intensity image is the brightness of pixels.
- They are dependent on a number of independent factors such as
  - object surface reflectance properties (given by the surface material, micro-structure and marking),
  - illumination properties,
  - object surface orientation with respect to a viewer and light source.
- This is a non-trivial and again ill-posed problem.

Ill-posed problem II

- The second problem is how to understand image brightness.
- The only information available in an intensity image is the brightness of pixels.
- They are dependent on a number of independent factors such as
  - object surface reflectance properties (given by the surface material, micro-structure and marking),
  - illumination properties,
  - object surface orientation with respect to a viewer and light source.
- This is a non-trivial and again ill-posed problem.

2D or 3D

- Some applications work with 2D images directly; for example,
  - an image of the flat specimen viewed by a microscope with transparent illumination;
  - a character drawn on a sheet of paper;
  - the image of a fingerprint, etc.
- Many basic and useful methods used in digital image analysis do not depend on whether the object was originally 2D or 3D (e.g, FFT).

2D or 3D

- Some applications work with 2D images directly; for example,
  - an image of the flat specimen viewed by a microscope with transparent illumination;
  - a character drawn on a sheet of paper;
  - the image of a fingerprint, etc.
- Many basic and useful methods used in digital image analysis do not depend on whether the object was originally 2D or 3D (e.g, FFT).

2D or 3D

- Some applications work with 2D images directly; for example,
  - an image of the flat specimen viewed by a microscope with transparent illumination;
  - a character drawn on a sheet of paper;
  - the image of a fingerprint, etc.
- Many basic and useful methods used in digital image analysis do not depend on whether the object was originally 2D or 3D (e.g, FFT).

2D or 3D

- Some applications work with 2D images directly; for example,
  - an image of the flat specimen viewed by a microscope with transparent illumination;
  - a character drawn on a sheet of paper;
  - the image of a fingerprint, etc.
- Many basic and useful methods used in digital image analysis do not depend on whether the object was originally 2D or 3D (e.g, FFT).

2D or 3D

- Some applications work with 2D images directly; for example,
  - an image of the flat specimen viewed by a microscope with transparent illumination;
  - a character drawn on a sheet of paper;
  - the image of a fingerprint, etc.
- Many basic and useful methods used in digital image analysis do not depend on whether the object was originally 2D or 3D (e.g, FFT).

Outline

Image formation: I

- The image formation process is described in [Horn, 1986, Wang and Wu, 1991].
- Related disciplines are photometry which is concerned with brightness measurement, and colorimetry which studies light reflectance or emission depending on wavelength.

Image formation: I

- The image formation process is described in [Horn, 1986, Wang and Wu, 1991].
- Related disciplines are photometry which is concerned with brightness measurement, and colorimetry which studies light reflectance or emission depending on wavelength.

Image formation: II

- A light source energy distribution $C(x, y, t, \lambda)$ depends in general on image co-ordinates (x, y), time t, and wavelength λ.
- For the human eye and most technical image sensors (e.g., TV cameras), the "brightness" f depends on the light source energy distribution C and the spectral sensitivity of the sensor, $S(\lambda)$ (dependent on the wavelength)

$$f(x, y, t) = \int_0^\infty C(x, y, t, \lambda) S(\lambda)\, d\lambda \tag{1}$$

- An intensity image $f(x, y, t)$ provides the brightness distribution.

Image formation: II

- A light source energy distribution $C(x, y, t, \lambda)$ depends in general on image co-ordinates (x, y), time t, and wavelength λ.
- For the human eye and most technical image sensors (e.g., TV cameras), the "brightness" f depends on the light source energy distribution C and the spectral sensitivity of the sensor, $S(\lambda)$ (dependent on the wavelength)

$$f(x, y, t) = \int_0^\infty C(x, y, t, \lambda) S(\lambda) \, d\lambda \qquad (1)$$

- An intensity image $f(x, y, t)$ provides the brightness distribution.

Image formation: II

- A light source energy distribution $C(x, y, t, \lambda)$ depends in general on image co-ordinates (x, y), time t, and wavelength λ.
- For the human eye and most technical image sensors (e.g., TV cameras), the “brightness” f depends on the light source energy distribution C and the spectral sensitivity of the sensor, $S(\lambda)$ (dependent on the wavelength)

$$f(x, y, t) = \int_0^\infty C(x, y, t, \lambda) S(\lambda)\, d\lambda \qquad (1)$$

- An intensity image $f(x, y, t)$ provides the brightness distribution.

Multi-spetral images

- In a color or multi-spectral image, the image is represented by a real vector function f

$$f(x, y, t) = (f_1(x, y, t), f_2(x, y, t), \cdots, f_n(x, y, t)) \quad (2)$$

where, for example, there may be red, green and blue, three components.

Monochromatic static image

- Image processing often deals with static images, in which time t is constant.
- A monochromatic static image is represented by a continuous image function $f(x, y)$ whose arguments are two co-ordinates in the plane.
- Most methods introduced in this course is primarily for intensity static image.
- It is often the case that the extension of the techniques to the multi-spectral case is obvious.

Monochromatic static image

- Image processing often deals with static images, in which time t is constant.
- A monochromatic static image is represented by a continuous image function $f(x, y)$ whose arguments are two co-ordinates in the plane.
- Most methods introduced in this course is primarily for intensity static image.
- It is often the case that the extension of the techniques to the multi-spectral case is obvious.

Monochromatic static image

- Image processing often deals with static images, in which time t is constant.
- A monochromatic static image is represented by a continuous image function $f(x, y)$ whose arguments are two co-ordinates in the plane.
- Most methods introduced in this course is primarily for intensity static image.
- It is often the case that the extension of the techniques to the multi-spectral case is obvious.

Monochromatic static image

- Image processing often deals with static images, in which time t is constant.
- A monochromatic static image is represented by a continuous image function $f(x, y)$ whose arguments are two co-ordinates in the plane.
- Most methods introduced in this course is primarily for intensity static image.
- It is often the case that the extension of the techniques to the multi-spectral case is obvious.

Monochromatic static image

- Computerized image processing uses digital image functions which are usually represented by matrices.
- Co-ordinates are integer numbers.
- The domain of the image function is a region R in the plane

$$R = \{(x, y) : 1 \leq x \leq x_m, 1 \leq y \leq y_n\} \tag{3}$$

where x_m and y_n represent maximal image co-ordinates.

Monochromatic static image

- Computerized image processing uses digital image functions which are usually represented by matrices.
- Co-ordinates are integer numbers.
- The domain of the image function is a region R in the plane

$$R = \{(x, y) : 1 \leq x \leq x_m, 1 \leq y \leq y_n\} \qquad (3)$$

where x_m and y_n represent maximal image co-ordinates.

Monochromatic static image

- Computerized image processing uses digital image functions which are usually represented by matrices.
- Co-ordinates are integer numbers.
- The domain of the image function is a region R in the plane

$$R = \{(x, y) : 1 \le x \le x_m, 1 \le y \le y_n\} \tag{3}$$

where x_m and y_n represent maximal image co-ordinates.

Limited domain

- The image function has a limited domain — infinite summation or integration limits can be used, as it is assumed that the image function is zero outside the domain.
- The customary orientation of co-ordinates in an image is in the normal Cartesian fashion (horizontal x axis, vertical y axis).
- The (row, column) orientation used in matrices is also quite often used in digital image processing.

Limited domain

- The image function has a limited domain — infinite summation or integration limits can be used, as it is assumed that the image function is zero outside the domain.
- The customary orientation of co-ordinates in an image is in the normal Cartesian fashion (horizontal x axis, vertical y axis).
- The (row, column) orientation used in matrices is also quite often used in digital image processing.

Limited domain

- The image function has a limited domain — infinite summation or integration limits can be used, as it is assumed that the image function is zero outside the domain.
- The customary orientation of co-ordinates in an image is in the normal Cartesian fashion (horizontal x axis, vertical y axis).
- The (row, column) orientation used in matrices is also quite often used in digital image processing.

Limited range

- The range of image function values is also limited; by convention, in intensity images the lowest value corresponds to black and the highest to white.
- Brightness values bounded by these limits are `gray levels`.
- The gray level range is $0, 1, \cdots, 255$, represented by 8 bits, the data type used is **unsigned char**. In some applications, 14 bits or more is used, e.g, for medical images.
- The usual computer display supports 8 bit gray level.

Limited range

- The range of image function values is also limited; by convention, in intensity images the lowest value corresponds to black and the highest to white.
- Brightness values bounded by these limits are `gray levels`.
- The gray level range is $0, 1, \cdots, 255$, represented by 8 bits, the data type used is **unsigned char**. In some applications, 14 bits or more is used, e.g, for medical images.
- The usual computer display supports 8 bit gray level.

Limited range

- The range of image function values is also limited; by convention, in intensity images the lowest value corresponds to black and the highest to white.
- Brightness values bounded by these limits are `gray levels`.
- The gray level range is $0, 1, \cdots, 255$, represented by 8 bits, the data type used is **unsigned char**. In some applications, 14 bits or more is used, e.g, for medical images.
- The usual computer display supports 8 bit gray level.

Limited range

- The range of image function values is also limited; by convention, in intensity images the lowest value corresponds to black and the highest to white.
- Brightness values bounded by these limits are `gray levels`.
- The gray level range is $0, 1, \cdots, 255$, represented by 8 bits, the data type used is **unsigned char**. In some applications, 14 bits or more is used, e.g, for medical images.
- The usual computer display supports 8 bit gray level.

Homework

- How to display a 16 bit gray level image? Generate an image of 16 bit and try to display it with your computer.
- If a discrete image is of continuous range, the image matrix is of type **float** or **double**. How to display it? Generate an image of **float** or **double** type and try to display it with your computer.

Homework

- How to display a 16 bit gray level image? Generate an image of 16 bit and try to display it with your computer.
- If a discrete image is of continuous range, the image matrix is of type **float** or **double**. How to display it? Generate an image of **float** or **double** type and try to display it with your computer.

Image quality

- The quality of a digital image grows in proportion to the spatial, spectral, radiometric, and time resolution.
- The `spatial resolution` is given by the proximity of image samples in the image plane.
- The `spectral resolution` is given by the bandwidth of the light frequencies captured by the sensor.
- The `radiometric resolution` (or `contrast resolution`, or `density resolution`) corresponds to the number of distinguishable gray levels.
- The `time resolution` is given by the interval between time samples at which images are captured.

Image quality

- The quality of a digital image grows in proportion to the spatial, spectral, radiometric, and time resolution.
- The `spatial resolution` is given by the proximity of image samples in the image plane.
- The `spectral resolution` is given by the bandwidth of the light frequencies captured by the sensor.
- The `radiometric resolution` (or `contrast resolution`, or `density resolution`) corresponds to the number of distinguishable gray levels.
- The `time resolution` is given by the interval between time samples at which images are captured.

Image quality

- The quality of a digital image grows in proportion to the spatial, spectral, radiometric, and time resolution.
- The `spatial resolution` is given by the proximity of image samples in the image plane.
- The `spectral resolution` is given by the bandwidth of the light frequencies captured by the sensor.
- The `radiometric resolution` (or `contrast resolution`, or `density resolution`) corresponds to the number of distinguishable gray levels.
- The `time resolution` is given by the interval between time samples at which images are captured.

Image quality

- The quality of a digital image grows in proportion to the spatial, spectral, radiometric, and time resolution.
- The `spatial resolution` is given by the proximity of image samples in the image plane.
- The `spectral resolution` is given by the bandwidth of the light frequencies captured by the sensor.
- The `radiometric resolution` (or `contrast resolution`, or `density resolution`) corresponds to the number of distinguishable gray levels.
- The `time resolution` is given by the interval between time samples at which images are captured.

Image quality

- The quality of a digital image grows in proportion to the spatial, spectral, radiometric, and time resolution.
- The `spatial resolution` is given by the proximity of image samples in the image plane.
- The `spectral resolution` is given by the bandwidth of the light frequencies captured by the sensor.
- The `radiometric resolution` (or `contrast resolution`, or `density resolution`) corresponds to the number of distinguishable gray levels.
- The `time resolution` is given by the interval between time samples at which images are captured.

Outline

Image as a stochastic process

- Images $f(x, y)$ can be treated as deterministic functions or as realizations of stochastic processes.
- Images are statistical in nature due to random changes and noise.
- It is sometimes of advantages to treat image functions as realizations of a stochastic process.

Image as a stochastic process

- Images $f(x, y)$ can be treated as deterministic functions or as realizations of stochastic processes.
- Images are statistical in nature due to random changes and noise.
- It is sometimes of advantages to treat image functions as realizations of a stochastic process.

Image as a stochastic process

- Images $f(x, y)$ can be treated as deterministic functions or as realizations of stochastic processes.
- Images are statistical in nature due to random changes and noise.
- It is sometimes of advantages to treat image functions as realizations of a stochastic process.

Image as a stochastic process: a typical model

- Mathematical tools used in image description have roots in linear system theory, integral transformations, discrete mathematics and the theory of stochastic processes, [Horn, 1986, Wang and Wu, 1991].
- A typical image formation model is described by a `linear spatial invariant system`,

$$f(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} g(a,b)h(x-a,y-b)\,dadb + n(x,y) \quad (4)$$

$$= h * g(x,y) + n(x,y) \quad (5)$$

where h is called the `point spread function` (PSF)and n is an additive noise.

Image as a stochastic process: a typical model

- Mathematical tools used in image description have roots in linear system theory, integral transformations, discrete mathematics and the theory of stochastic processes, [Horn, 1986, Wang and Wu, 1991].
- A typical image formation model is described by a `linear spatial invariant system`,

$$f(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} g(a,b)h(x-a,y-b)\,dadb + n(x,y) \quad (4)$$

$$= h * g(x,y) + n(x,y) \quad (5)$$

where h is called the `point spread function` (PSF)and n is an additive noise.

Image sampling and digitization

- An image captured by a sensor is expressed as a continuous function $f(x, y)$ of two co-ordinates in the plane.
- `Image digitization` means that the function $f(x, y)$ is sampled into a matrix with M rows and N columns.
- `Image quantization` assigns to each continuous sample an integer value. The continuous range of the image function $f(x, y)$ is split into K intervals.
- The finer the sampling (i.e., the larger M and N) and quantization (the larger K) the better the approximation of the continuous image function $f(x, y)$.

Image sampling and digitization

- An image captured by a sensor is expressed as a continuous function $f(x, y)$ of two co-ordinates in the plane.
- `Image digitization` means that the function $f(x, y)$ is sampled into a matrix with M rows and N columns.
- `Image quantization` assigns to each continuous sample an integer value. The continuous range of the image function $f(x, y)$ is split into K intervals.
- The finer the sampling (i.e., the larger M and N) and quantization (the larger K) the better the approximation of the continuous image function $f(x, y)$.

Image sampling and digitization

- An image captured by a sensor is expressed as a continuous function $f(x, y)$ of two co-ordinates in the plane.
- `Image digitization` means that the function $f(x, y)$ is sampled into a matrix with M rows and N columns.
- `Image quantization` assigns to each continuous sample an integer value. The continuous range of the image function $f(x, y)$ is split into K intervals.
- The finer the sampling (i.e., the larger M and N) and quantization (the larger K) the better the approximation of the continuous image function $f(x, y)$.

Image sampling and digitization

- An image captured by a sensor is expressed as a continuous function $f(x, y)$ of two co-ordinates in the plane.
- `Image digitization` means that the function $f(x, y)$ is sampled into a matrix with M rows and N columns.
- `Image quantization` assigns to each continuous sample an integer value. The continuous range of the image function $f(x, y)$ is split into K intervals.
- The finer the sampling (i.e., the larger M and N) and quantization (the larger K) the better the approximation of the continuous image function $f(x, y)$.

Two questions for sampling

- Two questions should be answered in connection with image function sampling:
  - First, the sampling period should be determined – the distance between two neighboring sampling points in the image;
  - Second, the geometric arrangement of sampling points (sampling grid) should be set.

Two questions for sampling

- Two questions should be answered in connection with image function sampling:
  - First, the sampling period should be determined – the distance between two neighboring sampling points in the image;
  - Second, the geometric arrangement of sampling points (sampling grid) should be set.

Two questions for sampling

- Two questions should be answered in connection with image function sampling:
  - First, the sampling period should be determined – the distance between two neighboring sampling points in the image;
  - Second, the geometric arrangement of sampling points (sampling grid) should be set.

Outline

Sampling intervals

- A continuous image function $f(x, y)$ can be sampled using a discrete grid of sampling points in the plane.
- The image is sampled at points $x = j\Delta x$, $y = k\Delta y$
- Two neighboring sampling points are separated by distances Δx along the x axis and Δy along the y axis.
- Distances Δx and Δy are called the sampling interval.
- The matrix of samples constitutes the discrete image.

Sampling intervals

- A continuous image function $f(x, y)$ can be sampled using a discrete grid of sampling points in the plane.
- The image is sampled at points $x = j\Delta x$, $y = k\Delta y$
- Two neighboring sampling points are separated by distances Δx along the x axis and Δy along the y axis.
- Distances Δx and Δy are called the sampling interval.
- The matrix of samples constitutes the discrete image.

Sampling intervals

- A continuous image function $f(x, y)$ can be sampled using a discrete grid of sampling points in the plane.
- The image is sampled at points $x = j\Delta x$, $y = k\Delta y$
- Two neighboring sampling points are separated by distances Δx along the x axis and Δy along the y axis.
- Distances Δx and Δy are called the sampling interval.
- The matrix of samples constitutes the discrete image.

Sampling intervals

- A continuous image function $f(x, y)$ can be sampled using a discrete grid of sampling points in the plane.
- The image is sampled at points $x = j\Delta x$, $y = k\Delta y$
- Two neighboring sampling points are separated by distances Δx along the x axis and Δy along the y axis.
- Distances Δx and Δy are called the sampling interval.
- The matrix of samples constitutes the discrete image.

Sampling intervals

- A continuous image function $f(x, y)$ can be sampled using a discrete grid of sampling points in the plane.
- The image is sampled at points $x = j\Delta x$, $y = k\Delta y$
- Two neighboring sampling points are separated by distances Δx along the x axis and Δy along the y axis.
- Distances Δx and Δy are called the sampling interval.
- The matrix of samples constitutes the discrete image.

Sampled Image

- The ideal sampling $s(x, y)$ in the regular grid can be represented using a collection of Dirac distributions

$$s(x, y) = \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \delta(x - j\Delta x, y - k\Delta y) \quad (6)$$

- The sampled image is the product of the continuous image $f(x, y)$ and the sampling function $s(x, y)$

$$f_s(x, y) = s(x, y) f(x, y) \quad (7)$$

Sampled Image

- The ideal sampling $s(x, y)$ in the regular grid can be represented using a collection of Dirac distributions

$$s(x, y) = \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \delta(x - j\Delta x, y - k\Delta y) \qquad (6)$$

- The sampled image is the product of the continuous image $f(x, y)$ and the sampling function $s(x, y)$

$$f_s(x, y) = s(x, y) f(x, y) \qquad (7)$$

Sampled Image

- The collection of Dirac distributions in equation (6) can be regarded as periodic with period Δx, Δy.
- It can be expanded into a Fourier series (assuming for a moment that the sampling grid covers the whole plane (infinite limits))

$$s = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} a_{mn} \mathbf{e}^{2\pi i \left(\frac{mx}{\Delta x} + \frac{ny}{\Delta y} \right)} \qquad (8)$$

Sampled Image

- The collection of Dirac distributions in equation (6) can be regarded as periodic with period Δx, Δy.
- It can be expanded into a Fourier series (assuming for a moment that the sampling grid covers the whole plane (infinite limits))

$$s = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} a_{mn} \mathbf{e}^{2\pi i \left(\frac{mx}{\Delta x} + \frac{ny}{\Delta y} \right)} \qquad (8)$$

Fourier expansion

- The coefficients of the Fourier expansion can be calculated as

$$a_{mn} = \frac{1}{\Delta x \Delta y} \int_{-\frac{\Delta x}{2}}^{\frac{\Delta x}{2}} \int_{-\frac{\Delta y}{2}}^{\frac{\Delta y}{2}} \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \delta(x - j\Delta x, y - k\Delta y) \mathbf{e}^{2\pi i \left(\frac{mx}{\Delta x} + \frac{ny}{\Delta y}\right)} dxdy \tag{9}$$

- Noting that only the term for $j = 0$ and $k = 0$ in the sum is nonzero in the range of integration (for other j and k, the center of the Delta function is outside the integral interval), the coefficients are

$$a_{mn} = \frac{1}{\Delta x \Delta y} \tag{10}$$

Fourier expansion

- The coefficients of the Fourier expansion can be calculated as

$$a_{mn} = \frac{1}{\Delta x \Delta y} \int_{-\frac{\Delta x}{2}}^{\frac{\Delta x}{2}} \int_{-\frac{\Delta y}{2}}^{\frac{\Delta y}{2}} \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \delta(x - j\Delta x, y - k\Delta y) \mathbf{e}^{2\pi i \left(\frac{mx}{\Delta x} + \frac{ny}{\Delta y} \right)} dxdy \tag{9}$$

- Noting that only the term for $j = 0$ and $k = 0$ in the sum is nonzero in the range of integration (for other j and k, the center of the Delta function is outside the integral interval), the coefficients are

$$a_{mn} = \frac{1}{\Delta x \Delta y} \tag{10}$$

Fourier expansion in the frequency domain

- Then, (7) can be rewritten as

$$f_s(x,y) = f(x,y)\frac{1}{\Delta x \Delta y}\sum_{m=-\infty}^{\infty}\sum_{n=-\infty}^{\infty} \mathbf{e}^{2\pi i\left(\frac{mx}{\Delta x}+\frac{ny}{\Delta y}\right)} \tag{11}$$

- In the frequency domain then

$$F_s(u,v) = \frac{1}{\Delta x \Delta y}\sum_{m=-\infty}^{\infty}\sum_{n=-\infty}^{\infty} F\left(u-\frac{m}{\Delta x}, v-\frac{n}{\Delta y}\right) \tag{12}$$

where F and F_s are the Fourier transform of f and f_s respectively.

Fourier expansion in the frequency domain

- Then, (7) can be rewritten as

$$f_s(x, y) = f(x, y)\frac{1}{\Delta x \Delta y}\sum_{m=-\infty}^{\infty}\sum_{n=-\infty}^{\infty} \mathbf{e}^{2\pi i\left(\frac{mx}{\Delta x}+\frac{ny}{\Delta y}\right)} \tag{11}$$

- In the frequency domain then

$$F_s(u, v) = \frac{1}{\Delta x \Delta y}\sum_{m=-\infty}^{\infty}\sum_{n=-\infty}^{\infty} F(u-\frac{m}{\Delta x}, v-\frac{n}{\Delta y}) \tag{12}$$

where F and F_s are the Fourier transform of f and f_s respectively.

Fourier transform of the sampled image

- Recall the Fourier transform is

$$F(u,v) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y)\mathbf{e}^{-2\pi i(ux+vy)}\,dxdy \qquad (13)$$

- Thus the Fourier transform of the sampled image is the sum of periodically repeated Fourier transforms $F(u,v)$ of the origin image.
- Periodic repetition of the Fourier transform result $F(u,v)$ may under certain conditions cause distortion of the image which is called `aliasing`.
- This happens when individual digitized components $F(u,v)$ overlap.

Fourier transform of the sampled image

- Recall the Fourier transform is

$$F(u, v) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \mathbf{e}^{-2\pi i (ux+vy)} \, dxdy \qquad (13)$$

- Thus the Fourier transform of the sampled image is the sum of periodically repeated Fourier transforms $F(u, v)$ of the origin image.
- Periodic repetition of the Fourier transform result $F(u, v)$ may under certain conditions cause distortion of the image which is called `aliasing`.
- This happens when individual digitized components $F(u, v)$ overlap.

Fourier transform of the sampled image

- Recall the Fourier transform is

$$F(u,v) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y)\mathbf{e}^{-2\pi i(ux+vy)}\,dxdy \qquad (13)$$

- Thus the Fourier transform of the sampled image is the sum of periodically repeated Fourier transforms $F(u,v)$ of the origin image.
- Periodic repetition of the Fourier transform result $F(u,v)$ may under certain conditions cause distortion of the image which is called `aliasing`.
- This happens when individual digitized components $F(u,v)$ overlap.

Fourier transform of the sampled image

- Recall the Fourier transform is

$$F(u,v)=\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y)\mathbf{e}^{-2\pi i(ux+vy)}\,dxdy \qquad (13)$$

- Thus the Fourier transform of the sampled image is the sum of periodically repeated Fourier transforms $F(u,v)$ of the origin image.
- Periodic repetition of the Fourier transform result $F(u,v)$ may under certain conditions cause distortion of the image which is called `aliasing`.
- This happens when individual digitized components $F(u,v)$ overlap.

Aliasing

- There is no aliasing if the image function $f(x, y)$ has a band limited spectrum, its Fourier transform $F(u, v) = 0$ outside a certain interval of frequencies $|u| > U$ and $|v| > V$.

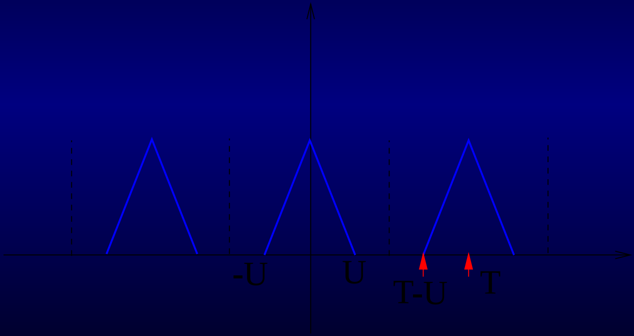

Figure: Where $T = \frac{1}{\Delta x}$.

Shannon sampling theorem

- From general sampling theory [Oppenheim et al., 1997], aliasing can be prevented if the sampling interval is chosen according to

$$\Delta x \leq \frac{1}{2U}, \qquad \Delta x \leq \frac{1}{2V}. \tag{14}$$

- This is the Shannon sampling theorem that has a simple physical interpretation in image analysis
  - the sampling interval should be chosen such that it is less than or equal to half of the smallest interesting detail in the image.

Shannon sampling theorem

- From general sampling theory [Oppenheim et al., 1997], aliasing can be prevented if the sampling interval is chosen according to

$$\Delta x \leq \frac{1}{2U}, \qquad \Delta x \leq \frac{1}{2V}. \tag{14}$$

- This is the Shannon sampling theorem that has a simple physical interpretation in image analysis
  - the sampling interval should be chosen such that it is less than or equal to half of the smallest interesting detail in the image.

Shannon sampling theorem

- From general sampling theory [Oppenheim et al., 1997], aliasing can be prevented if the sampling interval is chosen according to

$$\Delta x \leq \frac{1}{2U}, \qquad \Delta x \leq \frac{1}{2V}. \tag{14}$$

- This is the Shannon sampling theorem that has a simple physical interpretation in image analysis
  - the sampling interval should be chosen such that it is less than or equal to half of the smallest interesting detail in the image.

Sampling function in practice

- The sampling function is not the Dirac distribution in real digitizers – narrow impulses with limited amplitude are used instead.
- Assume a rectangular sampling grid which consists of $M \times N$ such equal and non-overlapping impulses $h_s(x, y)$ with sampling period Δx and Δy.
- Ideally, $h_s(x, y) = \delta(x, y)$.
- The function $h_s(x, y)$ simulates realistically the real image sensors.
- Outside the sensitive area of the sensor, the sampling element $h_s(x, y) = 0$.

Sampling function in practice

- The sampling function is not the Dirac distribution in real digitizers – narrow impulses with limited amplitude are used instead.
- Assume a rectangular sampling grid which consists of $M \times N$ such equal and non-overlapping impulses $h_s(x, y)$ with sampling period Δx and Δy.
- Ideally, $h_s(x, y) = \delta(x, y)$.
- The function $h_s(x, y)$ simulates realistically the real image sensors.
- Outside the sensitive area of the sensor, the sampling element $h_s(x, y) = 0$.

Sampling function in practice

- The sampling function is not the Dirac distribution in real digitizers – narrow impulses with limited amplitude are used instead.
- Assume a rectangular sampling grid which consists of $M \times N$ such equal and non-overlapping impulses $h_s(x, y)$ with sampling period Δx and Δy.
- Ideally, $h_s(x, y) = \delta(x, y)$.
- The function $h_s(x, y)$ simulates realistically the real image sensors.
- Outside the sensitive area of the sensor, the sampling element $h_s(x, y) = 0$.

Sampling function in practice

- The sampling function is not the Dirac distribution in real digitizers – narrow impulses with limited amplitude are used instead.
- Assume a rectangular sampling grid which consists of $M \times N$ such equal and non-overlapping impulses $h_s(x, y)$ with sampling period Δx and Δy.
- Ideally, $h_s(x, y) = \delta(x, y)$.
- The function $h_s(x, y)$ simulates realistically the real image sensors.
- Outside the sensitive area of the sensor, the sampling element $h_s(x, y) = 0$.

Sampling function in practice

- The sampling function is not the Dirac distribution in real digitizers – narrow impulses with limited amplitude are used instead.
- Assume a rectangular sampling grid which consists of $M \times N$ such equal and non-overlapping impulses $h_s(x, y)$ with sampling period Δx and Δy.
- Ideally, $h_s(x, y) = \delta(x, y)$.
- The function $h_s(x, y)$ simulates realistically the real image sensors.
- Outside the sensitive area of the sensor, the sampling element $h_s(x, y) = 0$.

Sampled image in practice

- The sampled image is then given by the following convolution

$$f_s(x,y) = f(x,y) \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} h_s(x - j\Delta x, y - k\Delta y) \quad (15)$$

- The sampled image f_s is distorted by the convolution of the original image f and the limited impulse h_s.
- The distortion of the frequency spectrum of the function F_s can be expressed as follows

$$F_s(u,v) = \frac{1}{\Delta x \Delta y} \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} F(u - \frac{m}{\Delta x}, v - \frac{n}{\Delta y}) H_s(\frac{m}{\Delta x}, \frac{n}{\Delta y}). \quad (16)$$

Sampled image in practice

- The sampled image is then given by the following convolution

$$f_s(x,y) = f(x,y) \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} h_s(x - j\Delta x, y - k\Delta y) \quad (15)$$

- The sampled image f_s is distorted by the convolution of the original image f and the limited impulse h_s.
- The distortion of the frequency spectrum of the function F_s can be expressed as follows

$$F_s(u,v) = \frac{1}{\Delta x \Delta y} \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} F(u - \frac{m}{\Delta x}, v - \frac{n}{\Delta y}) H_s(\frac{m}{\Delta x}, \frac{n}{\Delta y}). \quad (16)$$

Sampled image in practice

- The sampled image is then given by the following convolution

$$f_s(x, y) = f(x, y) \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} h_s(x - j\Delta x, y - k\Delta y) \tag{15}$$

- The sampled image f_s is distorted by the convolution of the original image f and the limited impulse h_s.
- The distortion of the frequency spectrum of the function F_s can be expressed as follows

$$F_s(u, v) = \frac{1}{\Delta x \Delta y} \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} F(u - \frac{m}{\Delta x}, v - \frac{n}{\Delta y}) H_s(\frac{m}{\Delta x}, \frac{n}{\Delta y}). \tag{16}$$

Homework

- Prove Eq. (16) from Eq. (15).

Homework

- There are other sampling schemes.
- These sampling points are ordered in the plane and their geometric relation is called the grid.
- Grids used in practice are mainly square or hexagonal

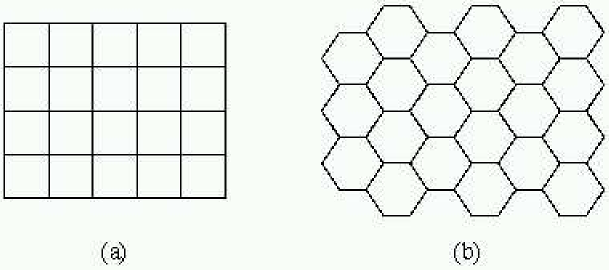

(a) Square grid, (b) hexagonal grid.

Homework

- There are other sampling schemes.
- These sampling points are ordered in the plane and their geometric relation is called the grid.
- Grids used in practice are mainly square or hexagonal

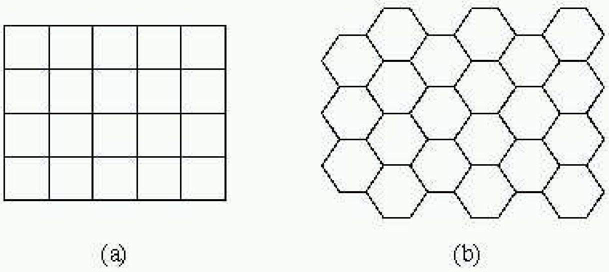

(a) Square grid, (b) hexagonal grid.

Homework

- There are other sampling schemes.
- These sampling points are ordered in the plane and their geometric relation is called the grid.
- Grids used in practice are mainly square or hexagonal

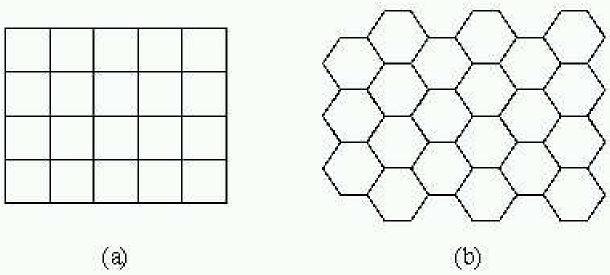

(a) Square grid, (b) hexagonal grid.

Pixels

- One infinitely small sampling point in the grid corresponds to one picture element (`pixel`) in the digital image.
- The set of pixels together covers the entire image.
- Pixels captured by a real digitization device have finite sizes.
- A pixel is a unit which is not further divisible.
- Sometimes pixels are also called points.

Pixels

- One infinitely small sampling point in the grid corresponds to one picture element (`pixel`) in the digital image.
- The set of pixels together covers the entire image.
- Pixels captured by a real digitization device have finite sizes.
- A pixel is a unit which is not further divisible.
- Sometimes pixels are also called points.

Pixels

- One infinitely small sampling point in the grid corresponds to one picture element (`pixel`) in the digital image.
- The set of pixels together covers the entire image.
- Pixels captured by a real digitization device have finite sizes.
- A pixel is a unit which is not further divisible.
- Sometimes pixels are also called points.

Pixels

- One infinitely small sampling point in the grid corresponds to one picture element (`pixel`) in the digital image.
- The set of pixels together covers the entire image.
- Pixels captured by a real digitization device have finite sizes.
- A pixel is a unit which is not further divisible.
- Sometimes pixels are also called points.

Pixels

- One infinitely small sampling point in the grid corresponds to one picture element (`pixel`) in the digital image.
- The set of pixels together covers the entire image.
- Pixels captured by a real digitization device have finite sizes.
- A pixel is a unit which is not further divisible.
- Sometimes pixels are also called points.

Real digitizers

- In real image digitizers, a sampling interval about ten times smaller than that indicated by the Shannon sampling theorem (14) is used.
- This is because algorithms which reconstruct the continuous image on a display from the digitized image function use only a step function.
- E.g., a line in the image is created from pixels represented by individual squares.
- The situation is more complicated.
- `Average sampling` happens in practice.

Real digitizers

- In real image digitizers, a sampling interval about ten times smaller than that indicated by the Shannon sampling theorem (14) is used.
- This is because algorithms which reconstruct the continuous image on a display from the digitized image function use only a step function.
- E.g., a line in the image is created from pixels represented by individual squares.

- The situation is more complicated.
- `Average sampling` happens in practice.

Real digitizers

- In real image digitizers, a sampling interval about ten times smaller than that indicated by the Shannon sampling theorem (14) is used.
- This is because algorithms which reconstruct the continuous image on a display from the digitized image function use only a step function.
- E.g., a line in the image is created from pixels represented by individual squares.

- The situation is more complicated.
- `Average sampling` happens in practice.

Real digitizers

- In real image digitizers, a sampling interval about ten times smaller than that indicated by the Shannon sampling theorem (14) is used.
- This is because algorithms which reconstruct the continuous image on a display from the digitized image function use only a step function.
- E.g., a line in the image is created from pixels represented by individual squares.
- The situation is more complicated.
- `Average sampling` happens in practice.

Real digitizers

- In real image digitizers, a sampling interval about ten times smaller than that indicated by the Shannon sampling theorem (14) is used.
- This is because algorithms which reconstruct the continuous image on a display from the digitized image function use only a step function.
- E.g., a line in the image is created from pixels represented by individual squares.

- The situation is more complicated.
- `Average sampling` happens in practice.

Sampling Examples

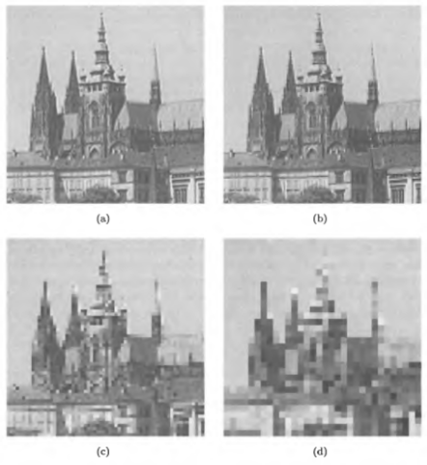

Figure: Digitizing, (a) 256 × 256. (b) 128 × 128. (c) 64 × 64. (d) 32 × 32. Images have been enlarged to the same size to illustrate the loss of detail.

Outline

Quantization

- The magnitude of a sampled image is expressed as a digital value in image processing.
- The transition between continuous values of the image function (brightness) and its digital equivalent is called `quantization`.
- The number of quantization levels should be high enough for human perception of fine shading details in the image.

Quantization

- The magnitude of a sampled image is expressed as a digital value in image processing.
- The transition between continuous values of the image function (brightness) and its digital equivalent is called `quantization`.
- The number of quantization levels should be high enough for human perception of fine shading details in the image.

Quantization

- The magnitude of a sampled image is expressed as a digital value in image processing.
- The transition between continuous values of the image function (brightness) and its digital equivalent is called `quantization`.
- The number of quantization levels should be high enough for human perception of fine shading details in the image.

False Contours in Images

- The occurrence of false contours is the main problem in image which have been quantized with insufficient brightness levels.
- This effect arises when the number of brightness levels is lower than that which humans can easily distinguish.

False Contours in Images

- The occurrence of false contours is the main problem in image which have been quantized with insufficient brightness levels.
- This effect arises when the number of brightness levels is lower than that which humans can easily distinguish.

False Contours in Color Images

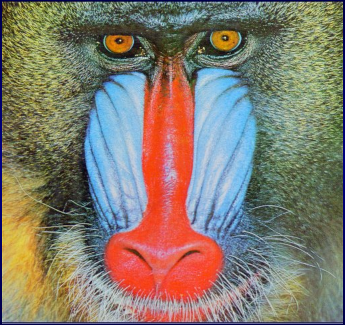

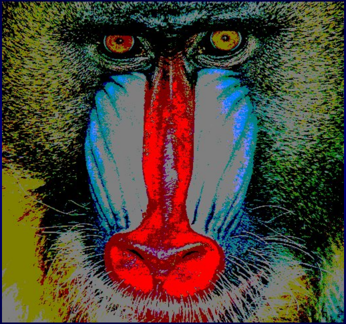

Discussions on Quantization I

- This number is dependent on many factors: e.g., the average local brightness.
- Displays which avoids this effect will normally provide a range of at least 100 intensity levels.
- This problem can be reduced when quantization into intervals of unequal length is used.
- The size of intervals corresponding to less probable brightnesses in the image is enlarged.
- These gray-scale transformation techniques are considered in later sections.

Discussions on Quantization II

- Most digital image processing devices use quantization into k equal intervals.
- If b bits are used ... the number of brightness levels is $k = 2^b$.
- Eight bits per pixel are commonly used, specialized measuring devices use 12 and more bits per pixel.

Quantization Experiment with matlab

Do you observe false contours when the quantization levels is decreasing?

Digital Properties

- A digital image has several properties,...
- both metric and topological,
- which are somewhat different from those of continuous two-dimensional functions we are familiar with.

Digital Properties

- A digital image has several properties,...
- both metric and topological,
- which are somewhat different from those of continuous two-dimensional functions we are familiar with.

Digital Properties

- A digital image has several properties,...
- both metric and topological,
- which are somewhat different from those of continuous two-dimensional functions we are familiar with.

Outline

Continuous Property May not Hold

- A digital image consists of picture elements of finite size.
- Usually pixels are arranged in a rectangular grid.
- A digital image is represented by a two-dimensional matrix whose elements are integer numbers corresponding to the quantization levels in the brightness scale.
- Some intuitively clear properties of continuous images have no straightforward analogy in the domain of digital images.

Continuous Property May not Hold

- A digital image consists of picture elements of finite size.
- Usually pixels are arranged in a rectangular grid.
- A digital image is represented by a two-dimensional matrix whose elements are integer numbers corresponding to the quantization levels in the brightness scale.
- Some intuitively clear properties of continuous images have no straightforward analogy in the domain of digital images.

Continuous Property May not Hold

- A digital image consists of picture elements of finite size.
- Usually pixels are arranged in a rectangular grid.
- A digital image is represented by a two-dimensional matrix whose elements are integer numbers corresponding to the quantization levels in the brightness scale.
- Some intuitively clear properties of continuous images have no straightforward analogy in the domain of digital images.

Continuous Property May not Hold

- A digital image consists of picture elements of finite size.
- Usually pixels are arranged in a rectangular grid.
- A digital image is represented by a two-dimensional matrix whose elements are integer numbers corresponding to the quantization levels in the brightness scale.
- Some intuitively clear properties of continuous images have no straightforward analogy in the domain of digital images.

Distance

- Distance is an important example.
- The distance between two pixels in a digital image is a significant quantitative measure.
- The distance between points with co-ordinates (i, j) and (h, k) may be defined in several different ways.

Distance

- Distance is an important example.
- The distance between two pixels in a digital image is a significant quantitative measure.
- The distance between points with co-ordinates (i, j) and (h, k) may be defined in several different ways.

Distance

- Distance is an important example.
- The distance between two pixels in a digital image is a significant quantitative measure.
- The distance between points with co-ordinates (i, j) and (h, k) may be defined in several different ways.

Euclidean Distance

- The Euclidean distance D_E is defined by

$$D_E[(i,j),h,k] = \sqrt{(i-h)^2 + (j-k)^2} \tag{17}$$

- The advantage of the Euclidean distance is the fact that it is intuitively obvious.
- The disadvantages are costly calculation due to the square root, and its not-integer value.

Euclidean Distance

- The Euclidean distance D_E is defined by

$$D_E[(i,j),h,k] = \sqrt{(i-h)^2+(j-k)^2} \qquad (17)$$

- The advantage of the Euclidean distance is the fact that it is intuitively obvious.
- The disadvantages are costly calculation due to the square root, and its not-integer value.

Euclidean Distance

- The Euclidean distance D_E is defined by

$$D_E[(i,j),h,k] = \sqrt{(i-h)^2 + (j-k)^2} \qquad (17)$$

- The advantage of the Euclidean distance is the fact that it is intuitively obvious.
- The disadvantages are costly calculation due to the square root, and its not-integer value.

TAXI Distance

- The distance between two points can also expressed as the minimum number of elementary steps in the digital grid which are needed to move from the starting point to the end point.
- If only horizontal and vertical moves are allowed, the distance D_4 or city block distance is obtained:

$$D_4[(i,j),h,k] = |i-h| + |j-k| \tag{18}$$

- This is the analogy with the distance between two locations in a city with a rectangular grid of streets and closed blocks of buildings.

TAXI Distance

- The distance between two points can also expressed as the minimum number of elementary steps in the digital grid which are needed to move from the starting point to the end point.
- If only horizontal and vertical moves are allowed, the distance D_4 or city block distance is obtained:

$$D_4[(i,j),h,k] = |i-h| + |j-k| \tag{18}$$

- This is the analogy with the distance between two locations in a city with a rectangular grid of streets and closed blocks of buildings.

TAXI Distance

- The distance between two points can also expressed as the minimum number of elementary steps in the digital grid which are needed to move from the starting point to the end point.
- If only horizontal and vertical moves are allowed, the distance D_4 or city block distance is obtained:

$$D_4[(i,j),h,k] = |i-h| + |j-k| \tag{18}$$

- This is the analogy with the distance between two locations in a city with a rectangular grid of streets and closed blocks of buildings.

Chess-board Distance

- If moves in diagonal directions are allowed in addition, the distance D_8 or the chess-board distance is obtained:

$$D_8[(i,j),h,k] = \max\{|i-h|,|j-k|\} \tag{19}$$

Pixel Neighborhoods

- Pixel adjacency is another important concept in digital images.
- Any two pixels are called 4-neighbors if they have distance $D_4 = 1$ from each other.
- 8-neighbors are two pixels with $D_8 = 1$.
- 4-neighbors and 8-neighbors are illustrated in Figure 7.

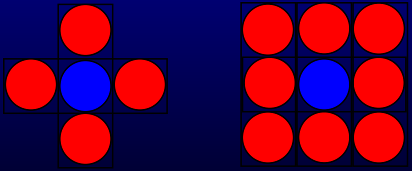

Figure: Pixel neighborhoods.

Pixel Neighborhoods

- Pixel adjacency is another important concept in digital images.
- Any two pixels are called 4-neighbors if they have distance $D_4 = 1$ from each other.
- 8-neighbors are two pixels with $D_8 = 1$.
- 4-neighbors and 8-neighbors are illustrated in Figure 7.

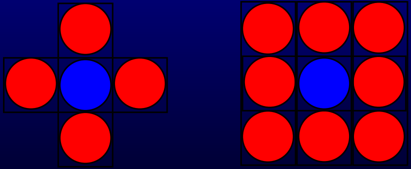

Figure: Pixel neighborhoods.

Pixel Neighborhoods

- Pixel adjacency is another important concept in digital images.
- Any two pixels are called 4-neighbors if they have distance $D_4 = 1$ from each other.
- 8-neighbors are two pixels with $D_8 = 1$.
- 4-neighbors and 8-neighbors are illustrated in Figure 7.

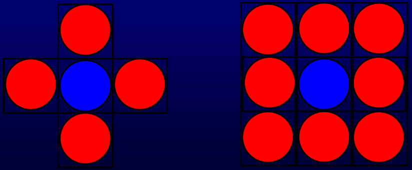

Figure: Pixel neighborhoods.

Pixel Neighborhoods

- Pixel adjacency is another important concept in digital images.
- Any two pixels are called 4-neighbors if they have distance $D_4 = 1$ from each other.
- 8-neighbors are two pixels with $D_8 = 1$.
- 4-neighbors and 8-neighbors are illustrated in Figure 7.

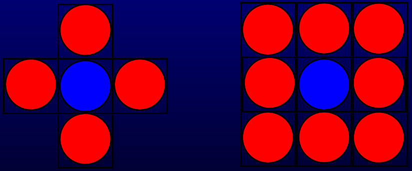

Figure: Pixel neighborhoods.

Regions and Paths

- It will become necessary to consider important sets consisting of several adjacent pixels — regions.
- `Region` is a contiguous (touching, neighboring, near to) set.
- A `path` from pixel P to pixel Q as a sequence of points $A_1, A_2, \cdots, A_n$, where $A_1 = P$ and $A_n = Q$, and A_{i+1} is a neighbor of A_i, $i = 1, \cdots, n-1$.
- A `region` is a set of pixels in which there is a path between any pair of its pixels, all of whose pixels also belong to the set.

Regions and Paths

- It will become necessary to consider important sets consisting of several adjacent pixels — regions.
- `Region` is a contiguous (touching, neighboring, near to) set.
- A `path` from pixel P to pixel Q as a sequence of points $A_1, A_2, \cdots, A_n$, where $A_1 = P$ and $A_n = Q$, and A_{i+1} is a neighbor of A_i, $i = 1, \cdots, n-1$.
- A `region` is a set of pixels in which there is a path between any pair of its pixels, all of whose pixels also belong to the set.

Regions and Paths

- It will become necessary to consider important sets consisting of several adjacent pixels — regions.
- `Region` is a contiguous (touching, neighboring, near to) set.
- A `path` from pixel P to pixel Q as a sequence of points $A_1, A_2, \cdots, A_n$, where $A_1 = P$ and $A_n = Q$, and A_{i+1} is a neighbor of A_i, $i = 1, \cdots, n-1$.
- A `region` is a set of pixels in which there is a path between any pair of its pixels, all of whose pixels also belong to the set.

Regions and Paths

- It will become necessary to consider important sets consisting of several adjacent pixels — regions.
- `Region` is a contiguous (touching, neighboring, near to) set.
- A `path` from pixel P to pixel Q as a sequence of points $A_1, A_2, \cdots, A_n$, where $A_1 = P$ and $A_n = Q$, and A_{i+1} is a neighbor of A_i, $i = 1, \cdots, n-1$.
- A `region` is a set of pixels in which there is a path between any pair of its pixels, all of whose pixels also belong to the set.

Contiguity

- If there is a path between two pixels in the set of pixels in the image, these pixels are called `contiguous`.
- The relation to be contiguous is reflexive, symmetric and transitive and therefore defines a decomposition of the set (in our case image) into equivalence classes (regions).
- The following image illustrates a binary image decomposed by the relation `contiguous` into three regions.

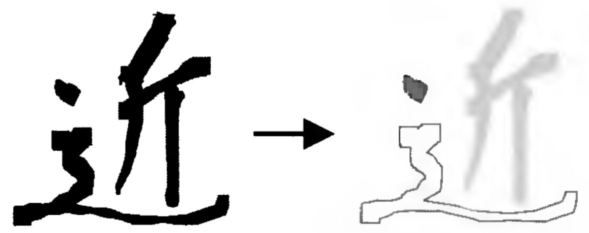

Contiguity

- If there is a path between two pixels in the set of pixels in the image, these pixels are called `contiguous`.
- The relation to be contiguous is reflexive, symmetric and transitive and therefore defines a decomposition of the set (in our case image) into equivalence classes (regions).
- The following image illustrates a binary image decomposed by the relation `contiguous` into three regions.

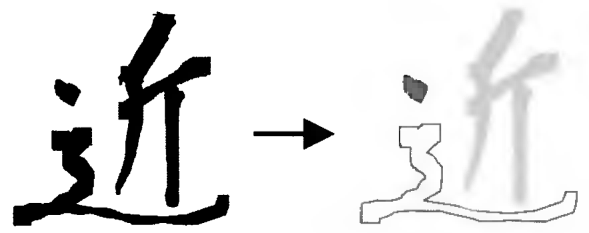

Contiguity

- If there is a path between two pixels in the set of pixels in the image, these pixels are called `contiguous`.
- The relation to be contiguous is reflexive, symmetric and transitive and therefore defines a decomposition of the set (in our case image) into equivalence classes (regions).
- The following image illustrates a binary image decomposed by the relation `contiguous` into three regions.

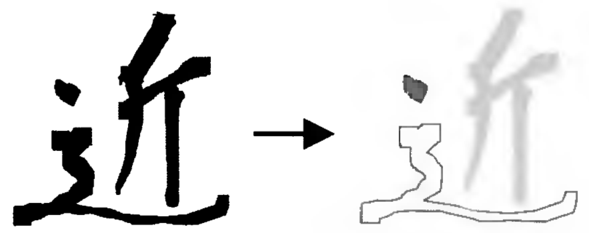

Some Terms

- Assume that R_i are disjoint regions in the image and that these regions do not touch the image boundary (to avoid special cases).
- Let R be the union of all regions R_i. Let R^C be the complement of R with respect to the image.
- The subset of R^C, which is contiguous with the image boundary, is called `background`, and the rest of the complement R^C is called `holes`.
- If there are no holes in a region we call it a `simply contiguous` region.
- A region with holes is called `multiply contiguous`.

Some Terms

- Assume that R_i are disjoint regions in the image and that these regions do not touch the image boundary (to avoid special cases).
- Let R be the union of all regions R_i. Let R^C be the complement of R with respect to the image.
- The subset of R^C, which is contiguous with the image boundary, is called `background`, and the rest of the complement R^C is called `holes`.
- If there are no holes in a region we call it a `simply contiguous` region.
- A region with holes is called `multiply contiguous`.

Some Terms

- Assume that R_i are disjoint regions in the image and that these regions do not touch the image boundary (to avoid special cases).
- Let R be the union of all regions R_i. Let R^C be the complement of R with respect to the image.
- The subset of R^C, which is contiguous with the image boundary, is called `background`, and the rest of the complement R^C is called `holes`.
- If there are no holes in a region we call it a `simply contiguous` region.
- A region with holes is called `multiply contiguous`.

Some Terms

- Assume that R_i are disjoint regions in the image and that these regions do not touch the image boundary (to avoid special cases).
- Let R be the union of all regions R_i. Let R^C be the complement of R with respect to the image.
- The subset of R^C, which is contiguous with the image boundary, is called `background`, and the rest of the complement R^C is called `holes`.
- If there are no holes in a region we call it a `simply contiguous` region.
- A region with holes is called `multiply contiguous`.

Some Terms

- Assume that R_i are disjoint regions in the image and that these regions do not touch the image boundary (to avoid special cases).
- Let R be the union of all regions R_i. Let R^C be the complement of R with respect to the image.
- The subset of R^C, which is contiguous with the image boundary, is called `background`, and the rest of the complement R^C is called `holes`.
- If there are no holes in a region we call it a `simply contiguous` region.
- A region with holes is called `multiply contiguous`.

Image Segmentation

- Note that the concept of region uses only the property to be contiguous.
- Secondary properties can be attached to regions which originate in image data interpretation.
- It is common to call some regions in the image `objects`.
- A process which determines which regions in an image correspond to objects in the world is part of `image segmentation`.

Image Segmentation

- Note that the concept of region uses only the property to be contiguous.
- Secondary properties can be attached to regions which originate in image data interpretation.
- It is common to call some regions in the image `objects`.
- A process which determines which regions in an image correspond to objects in the world is part of `image segmentation`.

Image Segmentation

- Note that the concept of region uses only the property to be contiguous.
- Secondary properties can be attached to regions which originate in image data interpretation.
- It is common to call some regions in the image `objects`.
- A process which determines which regions in an image correspond to objects in the world is part of `image segmentation`.

Image Segmentation

- Note that the concept of region uses only the property to be contiguous.
- Secondary properties can be attached to regions which originate in image data interpretation.
- It is common to call some regions in the image `objects`.
- A process which determines which regions in an image correspond to objects in the world is part of `image segmentation`.

Contiguous Example

- The brightness of a pixel is a property used to find objects in some images.
- If a pixel is darker than some other predefined values (threshold), then it belongs to some object.
- All such points which are also contiguous constitute one object.
- A hole consists of points which do not belong to the object and surrounded by the object, and all other points constitute the background.
- An example is the black printed text on the white paper, in which individual letters are objects.
- White areas surrounded by the letter are holes, e.g., the area inside a letter 'O'.
- Other parts of the paper are background.

Contiguous Example

- The brightness of a pixel is a property used to find objects in some images.
- If a pixel is darker than some other predefined values (threshold), then it belongs to some object.
- All such points which are also contiguous constitute one object.
- A hole consists of points which do not belong to the object and surrounded by the object, and all other points constitute the background.
- An example is the black printed text on the white paper, in which individual letters are objects.
- White areas surrounded by the letter are holes, e.g., the area inside a letter 'O'.
- Other parts of the paper are background.

Contiguous Example

- The brightness of a pixel is a property used to find objects in some images.
- If a pixel is darker than some other predefined values (threshold), then it belongs to some object.
- All such points which are also contiguous constitute one object.
- A hole consists of points which do not belong to the object and surrounded by the object, and all other points constitute the background.

- An example is the black printed text on the white paper, in which individual letters are objects.
- White areas surrounded by the letter are holes, e.g., the area inside a letter 'O'.
- Other parts of the paper are background.

Contiguous Example

- The brightness of a pixel is a property used to find objects in some images.
- If a pixel is darker than some other predefined values (threshold), then it belongs to some object.
- All such points which are also contiguous constitute one object.
- A hole consists of points which do not belong to the object and surrounded by the object, and all other points constitute the background.
- An example is the black printed text on the white paper, in which individual letters are objects.
- White areas surrounded by the letter are holes, e.g., the area inside a letter 'O'.
- Other parts of the paper are background.

Contiguous Example

- The brightness of a pixel is a property used to find objects in some images.
- If a pixel is darker than some other predefined values (threshold), then it belongs to some object.
- All such points which are also contiguous constitute one object.
- A hole consists of points which do not belong to the object and surrounded by the object, and all other points constitute the background.
- An example is the black printed text on the white paper, in which individual letters are objects.
- White areas surrounded by the letter are holes, e.g., the area inside a letter 'O'.
- Other parts of the paper are background.

Contiguous Example

- The brightness of a pixel is a property used to find objects in some images.
- If a pixel is darker than some other predefined values (threshold), then it belongs to some object.
- All such points which are also contiguous constitute one object.
- A hole consists of points which do not belong to the object and surrounded by the object, and all other points constitute the background.

- An example is the black printed text on the white paper, in which individual letters are objects.
- White areas surrounded by the letter are holes, e.g., the area inside a letter 'O'.
- Other parts of the paper are background.

Contiguous Example

- The brightness of a pixel is a property used to find objects in some images.
- If a pixel is darker than some other predefined values (threshold), then it belongs to some object.
- All such points which are also contiguous constitute one object.
- A hole consists of points which do not belong to the object and surrounded by the object, and all other points constitute the background.

- An example is the black printed text on the white paper, in which individual letters are objects.
- White areas surrounded by the letter are holes, e.g., the area inside a letter 'O'.
- Other parts of the paper are background.

Contiguity paradoxes

- These neighborhood and contiguity definitions on the square grid create some paradoxes.

Line Contiguous Paradox

The following figure shows three digital lines with 45^o and -45^o slope.

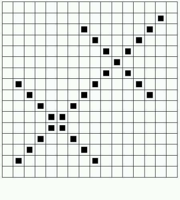

- If 4-connectivity is used, the lines are not contiguous at each of their points.
- An even worse conflict with intuitive understanding of line properties is:
  - two perpendicular lines do intersect in one case (upper right intersection) and do not intersect in another case (lower left), as they do not have any common point.

Line Contiguous Paradox

The following figure shows three digital lines with 45^o and -45^o slope.

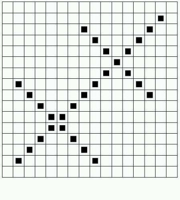

- If 4-connectivity is used, the lines are not contiguous at each of their points.
- An even worse conflict with intuitive understanding of line properties is:
  - two perpendicular lines do intersect in one case (upper right intersection) and do not intersect in another case (lower left), as they do not have any common point.

Line Contiguous Paradox

The following figure shows three digital lines with 45^o and -45^o slope.

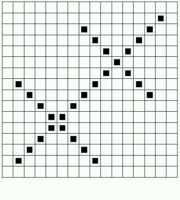

- If 4-connectivity is used, the lines are not contiguous at each of their points.
- An even worse conflict with intuitive understanding of line properties is:
  - two perpendicular lines do intersect in one case (upper right intersection) and do not intersect in another case (lower left), as they do not have any common point.

Jordon Paradox

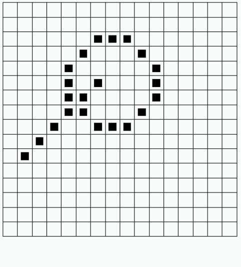

- Each closed curve divides the plane into two non-contiguous regions.
- If image are digitized in a square grid using 8-connectivity, there is a line from the inner part of a closed curve into the outer part without intersecting the curve.
- This implies that the inner and outer parts of the curve constitute one contiguous region.

Jordon Paradox

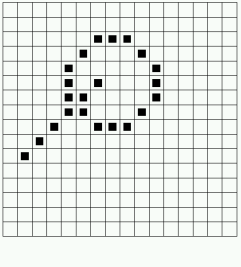

- Each closed curve divides the plane into two non-contiguous regions.
- If image are digitized in a square grid using 8-connectivity, there is a line from the inner part of a closed curve into the outer part without intersecting the curve.
- This implies that the inner and outer parts of the curve constitute one contiguous region.

Jordon Paradox

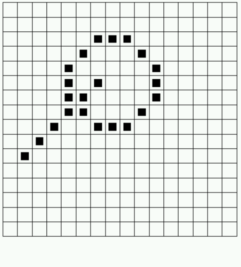

- Each closed curve divides the plane into two non-contiguous regions.
- If image are digitized in a square grid using 8-connectivity, there is a line from the inner part of a closed curve into the outer part without intersecting the curve.
- This implies that the inner and outer parts of the curve constitute one contiguous region.

Connectivity paradox

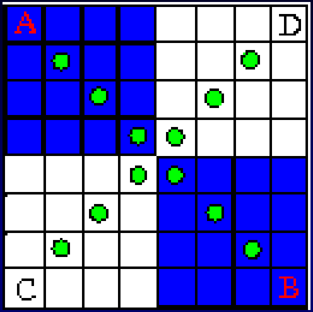

- If we assume 4-connectivity, the figure contains four separate contiguous regions *A*, *B*, *C* and *D*.
  - $A \cup B$ are disconnected, as well as $C \cup D$.
  - A topological contradiction.
  - Intuitively, $C \cup D$ should be connected if $A \cup B$ are disconnected.

Connectivity paradox

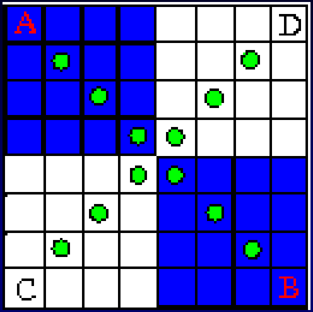

- If we assume 4-connectivity, the figure contains four separate contiguous regions A, B, C and D.
  - $A \cup B$ are disconnected, as well as $C \cup D$.
  - A topological contradiction.
  - Intuitively, $C \cup D$ should be connected if $A \cup B$ are disconnected.

Connectivity paradox

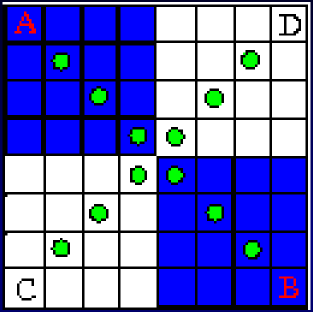

- If we assume 4-connectivity, the figure contains four separate contiguous regions *A*, *B*, *C* and *D*.
  - $A \cup B$ are disconnected, as well as $C \cup D$.
  - A topological contradiction.
  - Intuitively, $C \cup D$ should be connected if $A \cup B$ are disconnected.

Connectivity paradox

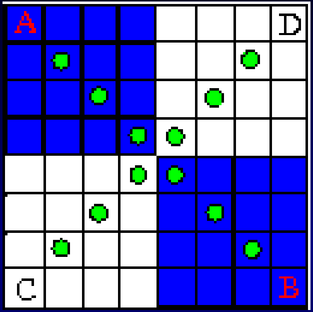

- If we assume 4-connectivity, the figure contains four separate contiguous regions *A*, *B*, *C* and *D*.
  - $A \cup B$ are disconnected, as well as $C \cup D$.
  - A topological contradiction.
  - Intuitively, $C \cup D$ should be connected if $A \cup B$ are disconnected.

Connectivity paradox

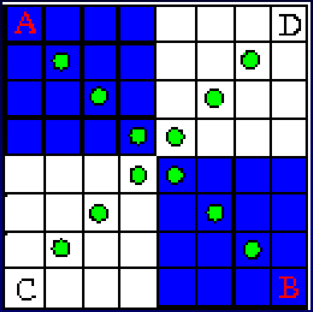

- If we assume 8-connectivity, there are two regions, $A \cup B$ and $C \cup D$.
- Both sets contain paths AB and CD entirely within themselves, but which also intersect!

Connectivity paradox

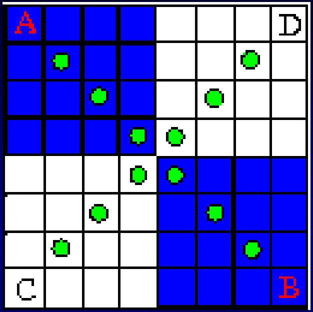

- If we assume 8-connectivity, there are two regions, $A \cup B$ and $C \cup D$.
- Both sets contain paths *AB* and *CD* entirely within themselves, but which also intersect!

Ad Hoc Solutions to Paradoxes

- One possible solution to contiguity paradox is to treat objects using 4-neighborhoods and background using 8-neighborhoods (or vice versa).
- More exact treatment of digital contiguity paradox and their solution for binary images and images with more brightness levels can be found in [Pavlidis, 1977].
- These problems are typical on square grids — a hexagonal grid (13) solves many of them.
- However, a grid of this type has also a number of disadvantages, [Pavlidis, 1977], p. 60.
- For reasons of simplicity and ease of processing, most digitizing devices use a square grid despite the stated drawbacks.
- We do not pursue further into this topic in this course, but use the simple approach, although there are some paradoxes.

Ad Hoc Solutions to Paradoxes

- One possible solution to contiguity paradox is to treat objects using 4-neighborhoods and background using 8-neighborhoods (or vice versa).
- More exact treatment of digital contiguity paradox and their solution for binary images and images with more brightness levels can be found in [Pavlidis, 1977].
- These problems are typical on square grids — a hexagonal grid (13) solves many of them.
- However, a grid of this type has also a number of disadvantages, [Pavlidis, 1977], p. 60.
- For reasons of simplicity and ease of processing, most digitizing devices use a square grid despite the stated drawbacks.
- We do not pursue further into this topic in this course, but use the simple approach, although there are some paradoxes.

Ad Hoc Solutions to Paradoxes

- One possible solution to contiguity paradox is to treat objects using 4-neighborhoods and background using 8-neighborhoods (or vice versa).
- More exact treatment of digital contiguity paradox and their solution for binary images and images with more brightness levels can be found in [Pavlidis, 1977].
- These problems are typical on square grids — a hexagonal grid (13) solves many of them.
- However, a grid of this type has also a number of disadvantages, [Pavlidis, 1977], p. 60.
- For reasons of simplicity and ease of processing, most digitizing devices use a square grid despite the stated drawbacks.
- We do not pursue further into this topic in this course, but use the simple approach, although there are some paradoxes.

Ad Hoc Solutions to Paradoxes

- One possible solution to contiguity paradox is to treat objects using 4-neighborhoods and background using 8-neighborhoods (or vice versa).
- More exact treatment of digital contiguity paradox and their solution for binary images and images with more brightness levels can be found in [Pavlidis, 1977].
- These problems are typical on square grids — a hexagonal grid (13) solves many of them.
- However, a grid of this type has also a number of disadvantages, [Pavlidis, 1977], p. 60.
- For reasons of simplicity and ease of processing, most digitizing devices use a square grid despite the stated drawbacks.
- We do not pursue further into this topic in this course, but use the simple approach, although there are some paradoxes.

Ad Hoc Solutions to Paradoxes

- One possible solution to contiguity paradox is to treat objects using 4-neighborhoods and background using 8-neighborhoods (or vice versa).
- More exact treatment of digital contiguity paradox and their solution for binary images and images with more brightness levels can be found in [Pavlidis, 1977].
- These problems are typical on square grids — a hexagonal grid (13) solves many of them.
- However, a grid of this type has also a number of disadvantages, [Pavlidis, 1977], p. 60.
- For reasons of simplicity and ease of processing, most digitizing devices use a square grid despite the stated drawbacks.
- We do not pursue further into this topic in this course, but use the simple approach, although there are some paradoxes.

Ad Hoc Solutions to Paradoxes

- One possible solution to contiguity paradox is to treat objects using 4-neighborhoods and background using 8-neighborhoods (or vice versa).
- More exact treatment of digital contiguity paradox and their solution for binary images and images with more brightness levels can be found in [Pavlidis, 1977].
- These problems are typical on square grids — a hexagonal grid (13) solves many of them.
- However, a grid of this type has also a number of disadvantages, [Pavlidis, 1977], p. 60.
- For reasons of simplicity and ease of processing, most digitizing devices use a square grid despite the stated drawbacks.
- We do not pursue further into this topic in this course, but use the simple approach, although there are some paradoxes.

Solutions to Paradoxes

- An alternative approach to the connectivity problems is to use discrete topology based on CW complex theory in topology.
- It is called cell complex in [Kovalevski, 1989].
- This approach develops a complete strand of image encoding and segmentation.
- The idea, first proposed by Riemann in the nineteenth century, considers families of sets of different dimensions:
  - points, which are 0-dimensional, may then be assigned to sets containing higher dimensional structures (such as pixel array), which permits the removal of the paradoxes we have seen.
  - line segments, which are 1-dimensional, gives precise definition of edge and border.

Solutions to Paradoxes

- An alternative approach to the connectivity problems is to use discrete topology based on CW complex theory in topology.
- It is called cell complex in [Kovalevski, 1989].
- This approach develops a complete strand of image encoding and segmentation.
- The idea, first proposed by Riemann in the nineteenth century, considers families of sets of different dimensions:
  - points, which are 0-dimensional, may then be assigned to sets containing higher dimensional structures (such as pixel array), which permits the removal of the paradoxes we have seen.
  - line segments, which are 1-dimensional, gives precise definition of edge and border.

Solutions to Paradoxes

- An alternative approach to the connectivity problems is to use discrete topology based on CW complex theory in topology.
- It is called cell complex in [Kovalevski, 1989].
- This approach develops a complete strand of image encoding and segmentation.
- The idea, first proposed by Riemann in the nineteenth century, considers families of sets of different dimensions:
  - points, which are 0-dimensional, may then be assigned to sets containing higher dimensional structures (such as pixel array), which permits the removal of the paradoxes we have seen.
  - line segments, which are 1-dimensional, gives precise definition of edge and border.

Solutions to Paradoxes

- An alternative approach to the connectivity problems is to use discrete topology based on CW complex theory in topology.
- It is called cell complex in [Kovalevski, 1989].
- This approach develops a complete strand of image encoding and segmentation.
- The idea, first proposed by Riemann in the nineteenth century, considers families of sets of different dimensions:
  - points, which are 0-dimensional, may then be assigned to sets containing higher dimensional structures (such as pixel array), which permits the removal of the paradoxes we have seen.
  - line segments, which are 1-dimensional, gives precise definition of edge and border.

Solutions to Paradoxes

- An alternative approach to the connectivity problems is to use discrete topology based on CW complex theory in topology.
- It is called cell complex in [Kovalevski, 1989].
- This approach develops a complete strand of image encoding and segmentation.
- The idea, first proposed by Riemann in the nineteenth century, considers families of sets of different dimensions:
  - points, which are 0-dimensional, may then be assigned to sets containing higher dimensional structures (such as pixel array), which permits the removal of the paradoxes we have seen.
  - line segments, which are 1-dimensional, gives precise definition of edge and border.

Solutions to Paradoxes

- An alternative approach to the connectivity problems is to use discrete topology based on CW complex theory in topology.
- It is called cell complex in [Kovalevski, 1989].
- This approach develops a complete strand of image encoding and segmentation.
- The idea, first proposed by Riemann in the nineteenth century, considers families of sets of different dimensions:
  - points, which are 0-dimensional, may then be assigned to sets containing higher dimensional structures (such as pixel array), which permits the removal of the paradoxes we have seen.
  - line segments, which are 1-dimensional, gives precise definition of edge and border.

Border

- Border of a region is another important concept in image analysis.
- The border of a region R is the set of pixels within the region that have one or more neighbors outside R.
- This definition of border is sometimes referred to as inner border, to distinguish it from the outer border,
  - it is the border of the background (i.e., the complement of) of the region.

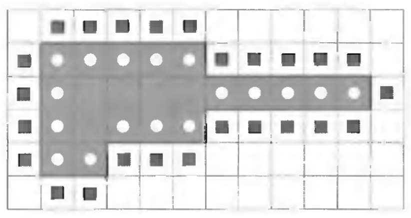

Border

- Border of a region is another important concept in image analysis.
- The border of a region R is the set of pixels within the region that have one or more neighbors outside R.
- This definition of border is sometimes referred to as inner border, to distinguish it from the outer border,
  - it is the border of the background (i.e., the complement of) of the region.

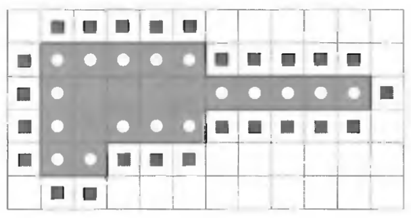

Border

- Border of a region is another important concept in image analysis.
- The border of a region R is the set of pixels within the region that have one or more neighbors outside R.
- This definition of border is sometimes referred to as inner border, to distinguish it from the outer border,
  - it is the border of the background (i.e., the complement of) of the region.

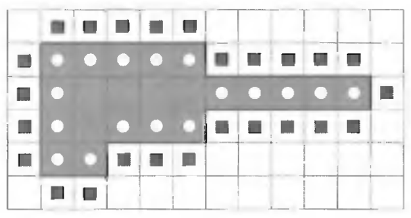

Border

- `Border` of a region is another important concept in image analysis.
- The border of a region R is the set of pixels within the region that have one or more neighbors outside R.
- This definition of border is sometimes referred to as `inner border`, to distinguish it from the `outer border`,
  - it is the border of the background (i.e., the complement of) of the region.

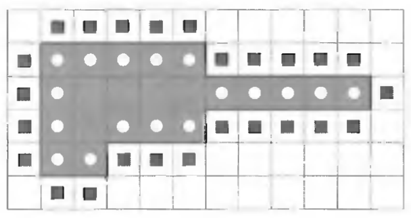

Edge

- Edge is a local property of a pixel and its immediate neighborhood — it is a vector given by a magnitude and direction.
- The edge direction is perpendicular to the gradient direction which points in the direction of image function growth.

Edge

- `Edge` is a local property of a pixel and its immediate neighborhood — it is a vector given by a magnitude and direction.
- The edge direction is perpendicular to the gradient direction which points in the direction of image function growth.

Border and Edge

- The border is a global concept related to a region, while edge expresses local properties of an image function.
- The border and edge are related as well.
- One possibility for finding boundaries is chaining the significant edges (points with high gradient of the image function).

Border and Edge

- The border is a global concept related to a region, while edge expresses local properties of an image function.
- The border and edge are related as well.
- One possibility for finding boundaries is chaining the significant edges (points with high gradient of the image function).

Border and Edge

- The border is a global concept related to a region, while edge expresses local properties of an image function.
- The border and edge are related as well.
- One possibility for finding boundaries is chaining the significant edges (points with high gradient of the image function).

Crack Edge I

- The edge property is attached to one pixel and its neighborhood.
- It is of advantage to assess properties between pairs of neighboring pixels.
- The concept of the `crack edge` comes from this idea.

Crack Edge I

- The edge property is attached to one pixel and its neighborhood.
- It is of advantage to assess properties between pairs of neighboring pixels.
- The concept of the `crack edge` comes from this idea.

Crack Edge I

- The edge property is attached to one pixel and its neighborhood.
- It is of advantage to assess properties between pairs of neighboring pixels.
- The concept of the `crack edge` comes from this idea.

Crack Edge II

- Four crack edges are attached to each pixel, which are defined by its relation to its 4-neighbors.
- The direction of the crack edge is that of increasing brightness, and is a multiple of 90 degrees.
- Its magnitude is the absolute difference between the brightness of the relevant pair of pixels.

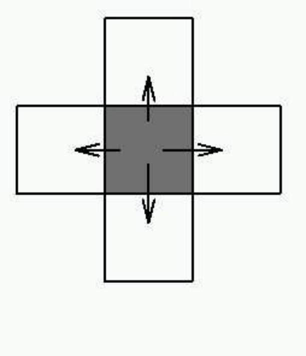

Crack Edge II

- Four crack edges are attached to each pixel, which are defined by its relation to its 4-neighbors.
- The direction of the crack edge is that of increasing brightness, and is a multiple of 90 degrees.
- Its magnitude is the absolute difference between the brightness of the relevant pair of pixels.

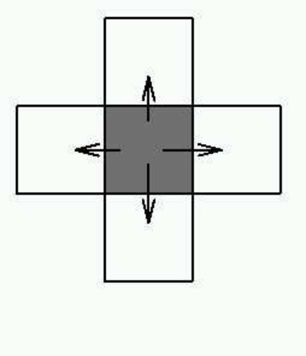

Crack Edge II

- Four crack edges are attached to each pixel, which are defined by its relation to its 4-neighbors.
- The direction of the crack edge is that of increasing brightness, and is a multiple of 90 degrees.
- Its magnitude is the absolute difference between the brightness of the relevant pair of pixels.

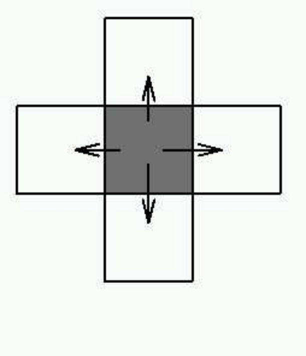

Convex Hull

- Convex hull is used to describe geometrical properties of objects.
- The convex hull is the smallest convex region which contains the object,
  - such that any two points of the region can be connected by a straight line, all points of which belong to the region.

Convex Hull

- Convex hull is used to describe geometrical properties of objects.
- The convex hull is the smallest convex region which contains the object,
  - such that any two points of the region can be connected by a straight line, all points of which belong to the region.

Convex Hull

- Convex hull is used to describe geometrical properties of objects.
- The convex hull is the smallest convex region which contains the object,
  - such that any two points of the region can be connected by a straight line, all points of which belong to the region.

Lakes and Bays

- An object can be represented by a collection of its topological components.
- The sets inside the convex hull which does not belong to an object is called the deficit of convexity.
- This can be split into two subsets.
  1. `lakes` are fully surrounded by the objects.
  2. `bays` are contiguous with the border of the convex hull of the object.
- The convex hull, lakes and bays are sometimes used for object description.

Lakes and Bays

- An object can be represented by a collection of its topological components.
- The sets inside the convex hull which does not belong to an object is called the deficit of convexity.
- This can be split into two subsets.
  1. `lakes` are fully surrounded by the objects.
  2. `bays` are contiguous with the border of the convex hull of the object.
- The convex hull, lakes and bays are sometimes used for object description.

Lakes and Bays

- An object can be represented by a collection of its topological components.
- The sets inside the convex hull which does not belong to an object is called the deficit of convexity.
- This can be split into two subsets.
  1. `lakes` are fully surrounded by the objects.
  2. `bays` are contiguous with the border of the convex hull of the object.
- The convex hull, lakes and bays are sometimes used for object description.

Lakes and Bays

- An object can be represented by a collection of its topological components.
- The sets inside the convex hull which does not belong to an object is called the deficit of convexity.
- This can be split into two subsets.
  1. lakes are fully surrounded by the objects.
  2. bays are contiguous with the border of the convex hull of the object.
- The convex hull, lakes and bays are sometimes used for object description.

Lakes and Bays

- An object can be represented by a collection of its topological components.
- The sets inside the convex hull which does not belong to an object is called the deficit of convexity.
- This can be split into two subsets.
  1. lakes are fully surrounded by the objects.
  2. bays are contiguous with the border of the convex hull of the object.
- The convex hull, lakes and bays are sometimes used for object description.

Lakes and Bays

- An object can be represented by a collection of its topological components.
- The sets inside the convex hull which does not belong to an object is called the deficit of convexity.
- This can be split into two subsets.
  1. `lakes` are fully surrounded by the objects.
  2. `bays` are contiguous with the border of the convex hull of the object.
- The convex hull, lakes and bays are sometimes used for object description.

Outline

Brightness Histogram

- Brightness histogram provides the frequency of the brightness value in the image.
- The brightness histogram $h_f(z)$ is a function showing,
  - for each brightness value z,
  - the number of pixels in the image f that have that brightness value z.

Brightness Histogram

- Brightness histogram provides the frequency of the brightness value in the image.
- The brightness histogram $h_f(z)$ is a function showing,
  - for each brightness value z,
  - the number of pixels in the image f that have that brightness value z.

Brightness Histogram

- Brightness histogram provides the frequency of the brightness value in the image.
- The brightness histogram $h_f(z)$ is a function showing,
  - for each brightness value z,
  - the number of pixels in the image f that have that brightness value z.

Brightness Histogram

- Brightness histogram provides the frequency of the brightness value in the image.
- The brightness histogram $h_f(z)$ is a function showing,
  - for each brightness value z,
  - the number of pixels in the image f that have that brightness value z.

Computing Brightness Histogram

- The histogram of an image with L gray levels is represented by a one-dimensional array with L elements.
- For a digital image of brightness value ranging in $[0, L-1]$, the following algorithm produces the brightness histogram:
  1. Assign zero values to all element of the array h_f;
  2. For all pixels (x, y) of the image f, increment $h_f[f(x, y)]$ by 1.

Computing Brightness Histogram

- The histogram of an image with L gray levels is represented by a one-dimensional array with L elements.
- For a digital image of brightness value ranging in $[0, L-1]$, the following algorithm produces the brightness histogram:
  1. Assign zero values to all element of the array h_f;
  2. For all pixels (x, y) of the image f, increment $h_f[f(x, y)]$ by 1.

Computing Brightness Histogram

- The histogram of an image with L gray levels is represented by a one-dimensional array with L elements.
- For a digital image of brightness value ranging in $[0, L-1]$, the following algorithm produces the brightness histogram:
  1. Assign zero values to all element of the array h_f;
  2. For all pixels (x, y) of the image f, increment $h_f[f(x, y)]$ by 1.

Computing Brightness Histogram

- The histogram of an image with L gray levels is represented by a one-dimensional array with L elements.
- For a digital image of brightness value ranging in $[0, L-1]$, the following algorithm produces the brightness histogram:
  1. Assign zero values to all element of the array h_f;
  2. For all pixels (x, y) of the image f, increment $h_f[f(x, y)]$ by 1.

Brightness Histogram Example

- The histogram is often displayed as a bar graph.

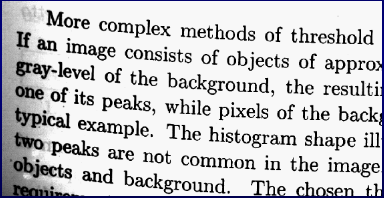

More complex methods of threshold
If an image consists of objects of approx
gray-level of the background, the resulti
one of its peaks, while pixels of the back
typical example. The histogram shape ill
two peaks are not common in the image
objects and background. The chosen th

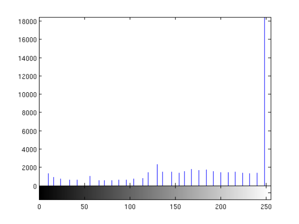

Computation Options

- Computing brightness histogram is similar to generating the histogram of a random variable from a given group of samples.
- In the above algorithm, the starting value is 0, bin-width 1, and bin number *L*.
- This algorithm can be modified to generate brightness histogram of arbitrary bin-width and bin number.
- For multi-spectral band images, histogram of each individual band can be generated in a similar way.

Computation Options

- Computing brightness histogram is similar to generating the histogram of a random variable from a given group of samples.
- In the above algorithm, the starting value is 0, bin-width 1, and bin number L.
- This algorithm can be modified to generate brightness histogram of arbitrary bin-width and bin number.
- For multi-spectral band images, histogram of each individual band can be generated in a similar way.

Computation Options

- Computing brightness histogram is similar to generating the histogram of a random variable from a given group of samples.
- In the above algorithm, the starting value is 0, bin-width 1, and bin number L.
- This algorithm can be modified to generate brightness histogram of arbitrary bin-width and bin number.
- For multi-spectral band images, histogram of each individual band can be generated in a similar way.

Computation Options

- Computing brightness histogram is similar to generating the histogram of a random variable from a given group of samples.
- In the above algorithm, the starting value is 0, bin-width 1, and bin number L.
- This algorithm can be modified to generate brightness histogram of arbitrary bin-width and bin number.
- For multi-spectral band images, histogram of each individual band can be generated in a similar way.

Brightness Histogram Example: non-convential bin

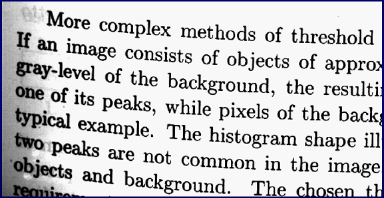
More complex methods of threshold
If an image consists of objects of approx
gray-level of the background, the resulti
one of its peaks, while pixels of the back
typical example. The histogram shape ill
two peaks are not common in the image
objects and background. The chosen th

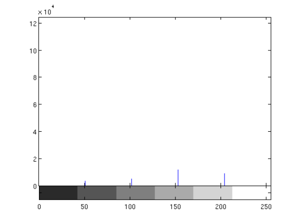

Histogram: discussions I

- The histogram provides a natural bridge between images and a probabilistic description.
- Wc might want to find a first-order probability function $p_I(z; x, y)$ to indicate the probability that pixel (x, y) has brightness z.
- Dependence on the position of the pixel is not of interest in the histogram.

Histogram: discussions I

- The histogram provides a natural bridge between images and a probabilistic description.
- Wc might want to find a first-order probability function $p_I(z; x, y)$ to indicate the probability that pixel (x, y) has brightness z.
- Dependence on the position of the pixel is not of interest in the histogram.

Histogram: discussions I

- The histogram provides a natural bridge between images and a probabilistic description.
- Wc might want to find a first-order probability function $p_I(z; x, y)$ to indicate the probability that pixel (x, y) has brightness z.
- Dependence on the position of the pixel is not of interest in the histogram.

Histogram: discussions II

- The histogram is usually the only global information about the image which is available.
- It is used when finding optimal illumination conditions for capturing an image, gray-scale transformations, and image segmentation to objects and background.
- Note that one histogram may correspond to several images;
  - e.g., a change of the object position on a constant background does not affect the histogram.

Histogram: discussions II

- The histogram is usually the only global information about the image which is available.
- It is used when finding optimal illumination conditions for capturing an image, gray-scale transformations, and image segmentation to objects and background.
- Note that one histogram may correspond to several images;
  - e.g., a change of the object position on a constant background does not affect the histogram.

Histogram: discussions II

- The histogram is usually the only global information about the image which is available.
- It is used when finding optimal illumination conditions for capturing an image, gray-scale transformations, and image segmentation to objects and background.
- Note that one histogram may correspond to several images;
  - e.g., a change of the object position on a constant background does not affect the histogram.

Histogram: discussions II

- The histogram is usually the only global information about the image which is available.
- It is used when finding optimal illumination conditions for capturing an image, gray-scale transformations, and image segmentation to objects and background.
- Note that one histogram may correspond to several images;
  - e.g., a change of the object position on a constant background does not affect the histogram.

Histogram: discussions II

- The histogram of a digital image typically has many local minima and maxima, which may complicate its further processing.
- This problem can be avoided by local smoothing of the histogram.
- This algorithm would need some boundary adjustment, and carries no guarantee of removing all local minima.
- Other techniques for smoothing exist, notably Gaussian blurring.

Histogram: discussions II

- The histogram of a digital image typically has many local minima and maxima, which may complicate its further processing.
- This problem can be avoided by local smoothing of the histogram.
- This algorithm would need some boundary adjustment, and carries no guarantee of removing all local minima.
- Other techniques for smoothing exist, notably Gaussian blurring.

Histogram: discussions II

- The histogram of a digital image typically has many local minima and maxima, which may complicate its further processing.
- This problem can be avoided by local smoothing of the histogram.
- This algorithm would need some boundary adjustment, and carries no guarantee of removing all local minima.
- Other techniques for smoothing exist, notably Gaussian blurring.

Histogram: discussions II

- The histogram of a digital image typically has many local minima and maxima, which may complicate its further processing.
- This problem can be avoided by local smoothing of the histogram.
- This algorithm would need some boundary adjustment, and carries no guarantee of removing all local minima.
- Other techniques for smoothing exist, notably Gaussian blurring.

Histogram Experiment with matlab

- Creating an image histogram using imhist.

Outline

Entropy

- Image information content can be estimated using entropy H.
- The concept of `entropy` has roots in thermodynamics and statistical mechanics, but it took many years before entropy was related to information.
- The information-theoretic formulation of entropy comes from Shannon [Shannon, 1948] and is often called `information entropy`.

Entropy

- Image information content can be estimated using entropy H.
- The concept of `entropy` has roots in thermodynamics and statistical mechanics, but it took many years before entropy was related to information.
- The information-theoretic formulation of entropy comes from Shannon [Shannon, 1948] and is often called `information entropy`.

Entropy

- Image information content can be estimated using entropy *H*.
- The concept of `entropy` has roots in thermodynamics and statistical mechanics, but it took many years before entropy was related to information.
- The information-theoretic formulation of entropy comes from Shannon [Shannon, 1948] and is often called `information entropy`.

Entropy As a Measure of uncertainty

- An intuitive understanding of information entropy relates to the amount of uncertainty about an event associated with a given probability distribution.
- The entropy can serve as an measure of “disorder”.
- As the level of disorder rises, entropy increases and events are less predictable.

Entropy As a Measure of uncertainty

- An intuitive understanding of information entropy relates to the amount of uncertainty about an event associated with a given probability distribution.
- The entropy can serve as an measure of “disorder”.
- As the level of disorder rises, entropy increases and events are less predictable.

Entropy As a Measure of uncertainty

- An intuitive understanding of information entropy relates to the amount of uncertainty about an event associated with a given probability distribution.
- The entropy can serve as an measure of “disorder”.
- As the level of disorder rises, entropy increases and events are less predictable.

Entropy: definition

- The entropy is defined formally assuming a discrete random variable X with possible outcomes (called also states) $x_1, \cdots, x_n$.
- Let $p(x_k)$ be the probability of the outcome x_k, $k = 1, \cdots, n$.
- Then the entropy is defined as

$$H(x) = \sum_{k=1}^{n} p(x_k) \log \frac{1}{p(x_k)} = -\sum_{k=1}^{n} p(x_k) \log p(x_k). \tag{20}$$

- $\log \frac{1}{p(x_k)}$ is called the surprisal of the outcome x_k.
- The entropy is the expected value of its outcome's surprisal.

Entropy: definition

- The entropy is defined formally assuming a discrete random variable X with possible outcomes (called also states) $x_1, \cdots, x_n$.
- Let $p(x_k)$ be the probability of the outcome x_k, $k = 1, \cdots, n$.
- Then the entropy is defined as

$$H(x) = \sum_{k=1}^{n} p(x_k) \log \frac{1}{p(x_k)} = -\sum_{k=1}^{n} p(x_k) \log p(x_k). \tag{20}$$

- $\log \frac{1}{p(x_k)}$ is called the surprisal of the outcome x_k.
- The entropy is the expected value of its outcome's surprisal.

Entropy: definition

- The entropy is defined formally assuming a discrete random variable X with possible outcomes (called also states) $x_1, \cdots, x_n$.
- Let $p(x_k)$ be the probability of the outcome x_k, $k = 1, \cdots, n$.
- Then the entropy is defined as

$$H(x) = \sum_{k=1}^{n} p(x_k) \log \frac{1}{p(x_k)} = -\sum_{k=1}^{n} p(x_k) \log p(x_k). \tag{20}$$

- $\log \frac{1}{p(x_k)}$ is called the surprisal of the outcome x_k.
- The entropy is the expected value of its outcome's surprisal.

Entropy: definition

- The entropy is defined formally assuming a discrete random variable X with possible outcomes (called also states) $x_1, \cdots, x_n$.
- Let $p(x_k)$ be the probability of the outcome x_k, $k = 1, \cdots, n$.
- Then the entropy is defined as

$$H(x) = \sum_{k=1}^{n} p(x_k) \log \frac{1}{p(x_k)} = -\sum_{k=1}^{n} p(x_k) \log p(x_k). \tag{20}$$

- $\log \frac{1}{p(x_k)}$ is called the surprisal of the outcome x_k.
- The entropy is the expected value of its outcome's surprisal.

Entropy: definition

- The entropy is defined formally assuming a discrete random variable X with possible outcomes (called also states) $x_1, \cdots, x_n$.
- Let $p(x_k)$ be the probability of the outcome x_k, $k = 1, \cdots, n$.
- Then the entropy is defined as

$$H(x) = \sum_{k=1}^{n} p(x_k) \log \frac{1}{p(x_k)} = -\sum_{k=1}^{n} p(x_k) \log p(x_k). \tag{20}$$

- $\log \frac{1}{p(x_k)}$ is called the surprisal of the outcome x_k.
- The entropy is the expected value of its outcome's surprisal.

Entropy: discussions

- The base of the logarithm in this formula determines the unit in which entropy is measured.
- If this base is two then the entropy is given in bits.
- The entropy is often estimated using a gray-level histogram in image analysis.
- Because entropy measures the uncertainty about the realization of a random variable, it is used to assess redundancy in an image for image compression.

Entropy: discussions

- The base of the logarithm in this formula determines the unit in which entropy is measured.
- If this base is two then the entropy is given in bits.
- The entropy is often estimated using a gray-level histogram in image analysis.
- Because entropy measures the uncertainty about the realization of a random variable, it is used to assess redundancy in an image for image compression.

Entropy: discussions

- The base of the logarithm in this formula determines the unit in which entropy is measured.
- If this base is two then the entropy is given in bits.
- The entropy is often estimated using a gray-level histogram in image analysis.
- Because entropy measures the uncertainty about the realization of a random variable, it is used to assess redundancy in an image for image compression.

Entropy: discussions

- The base of the logarithm in this formula determines the unit in which entropy is measured.
- If this base is two then the entropy is given in bits.
- The entropy is often estimated using a gray-level histogram in image analysis.
- Because entropy measures the uncertainty about the realization of a random variable, it is used to assess redundancy in an image for image compression.

Outline

Basic concepts

- Image functions
- Image formation
- Image as a stochastic process

Image digitization

- Sampling
- Quantization

Digital image properties

- Metric and topological properties of digital images
  - Metric properties of digital images
  - Topological properties of digital images
  - Contiguity paradoxes
  - Other topological and geometrical properties
- Histogram
- Entropy
- Visual perception of the image
  - Contrast
  - Acuity
  - Visual Illusions
- Image quality

Visual Perception

- Anyone who creates or uses algorithms or devices for digital image processing should take into account the principle of human visual perception.
- There are psycho-physical parameters such as contrast, border, shape, texture, color, etc.
- Humans will find objects in images only if they may be distinguished effortlessly from the background.
- Human perception of image provokes many illusions, the understanding of which provides valuable clues about visual mechanisms.
- The topic is covered exhaustively from the point of view of computer vision in [Frisby, 1979].

Visual Perception

- Anyone who creates or uses algorithms or devices for digital image processing should take into account the principle of human visual perception.
- There are psycho-physical parameters such as contrast, border, shape, texture, color, etc.
- Humans will find objects in images only if they may be distinguished effortlessly from the background.
- Human perception of image provokes many illusions, the understanding of which provides valuable clues about visual mechanisms.
- The topic is covered exhaustively from the point of view of computer vision in [Frisby, 1979].

Visual Perception

- Anyone who creates or uses algorithms or devices for digital image processing should take into account the principle of human visual perception.
- There are psycho-physical parameters such as contrast, border, shape, texture, color, etc.
- Humans will find objects in images only if they may be distinguished effortlessly from the background.
- Human perception of image provokes many illusions, the understanding of which provides valuable clues about visual mechanisms.
- The topic is covered exhaustively from the point of view of computer vision in [Frisby, 1979].

Visual Perception

- Anyone who creates or uses algorithms or devices for digital image processing should take into account the principle of human visual perception.
- There are psycho-physical parameters such as contrast, border, shape, texture, color, etc.
- Humans will find objects in images only if they may be distinguished effortlessly from the background.
- Human perception of image provokes many illusions, the understanding of which provides valuable clues about visual mechanisms.
- The topic is covered exhaustively from the point of view of computer vision in [Frisby, 1979].

Visual Perception

- Anyone who creates or uses algorithms or devices for digital image processing should take into account the principle of human visual perception.
- There are psycho-physical parameters such as contrast, border, shape, texture, color, etc.
- Humans will find objects in images only if they may be distinguished effortlessly from the background.
- Human perception of image provokes many illusions, the understanding of which provides valuable clues about visual mechanisms.
- The topic is covered exhaustively from the point of view of computer vision in [Frisby, 1979].

Contrast

- Contrast is the local change in brightness and is defined as the ratio between average brightness of an object and the background brightness.
- The human eye is logarithmically sensitive to brightness.
- Gamma correction is used to calibrate the differences among different computer monitors.
- Apparent brightness depends very much on the brightness of the local background; this effect is called conditional contrast.

Contrast

- Contrast is the local change in brightness and is defined as the ratio between average brightness of an object and the background brightness.
- The human eye is logarithmically sensitive to brightness.
- Gamma correction is used to calibrate the differences among different computer monitors.
- Apparent brightness depends very much on the brightness of the local background; this effect is called conditional contrast.

Contrast

- Contrast is the local change in brightness and is defined as the ratio between average brightness of an object and the background brightness.
- The human eye is logarithmically sensitive to brightness.
- Gamma correction is used to calibrate the differences among different computer monitors.
- Apparent brightness depends very much on the brightness of the local background; this effect is called conditional contrast.

Contrast

- Contrast is the local change in brightness and is defined as the ratio between average brightness of an object and the background brightness.
- The human eye is logarithmically sensitive to brightness.
- Gamma correction is used to calibrate the differences among different computer monitors.
- Apparent brightness depends very much on the brightness of the local background; this effect is called conditional contrast.

Conditional Contrast

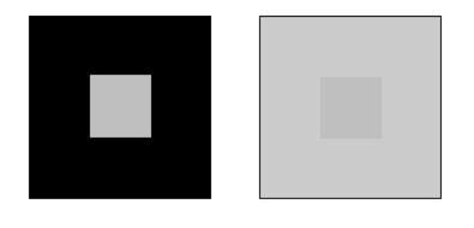

- The figure illustrates this effect with two small squares of the same brightness on a dark and a light background.
- Human perceives the brightness of the small squares as different.

Conditional Contrast

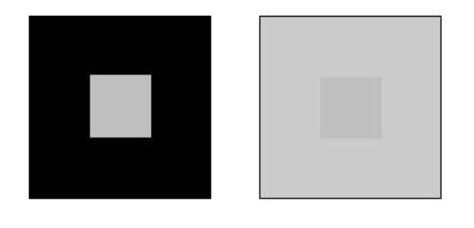

- The figure illustrates this effect with two small squares of the same brightness on a dark and a light background.
- Human perceives the brightness of the small squares as different.

Acuity

- Acuity is the ability to detect details in image.
- The human eye is less sensitive to slow and fast brightness changes but is more sensitive to intermediate changes.
- Resolution in an image is firmly bounded by the resolution ability of the human eye;
- there is no sense in representing visual information with higher resolution than that of the viewer.

Acuity

- Acuity is the ability to detect details in image.
- The human eye is less sensitive to slow and fast brightness changes but is more sensitive to intermediate changes.
- Resolution in an image is firmly bounded by the resolution ability of the human eye;
- there is no sense in representing visual information with higher resolution than that of the viewer.

Acuity

- Acuity is the ability to detect details in image.
- The human eye is less sensitive to slow and fast brightness changes but is more sensitive to intermediate changes.
- Resolution in an image is firmly bounded by the resolution ability of the human eye;
- there is no sense in representing visual information with higher resolution than that of the viewer.

Acuity

- Acuity is the ability to detect details in image.
- The human eye is less sensitive to slow and fast brightness changes but is more sensitive to intermediate changes.
- Resolution in an image is firmly bounded by the resolution ability of the human eye;
- there is no sense in representing visual information with higher resolution than that of the viewer.

Resolution

- Resolution in optics is defined as the inverse value of a maximum viewing angle between the viewer and two proximate points which human cannot distinguish, and so fuse together.
- Human vision has the best resolution for objects which are at a distance of about 25 cm from an eye under illumination of about 500 lux.
- This illumination is provided by a 60 W from a distance 40 cm.
- Under this conditions the distance between two distinguishable points is approximately 0.16 mm.
- Another report says that the minimal distinguishable distance is 0.47 mm [Kutter, 1999].

Resolution

- Resolution in optics is defined as the inverse value of a maximum viewing angle between the viewer and two proximate points which human cannot distinguish, and so fuse together.
- Human vision has the best resolution for objects which are at a distance of about 25 cm from an eye under illumination of about 500 lux.
- This illumination is provided by a 60 W from a distance 40 cm.
- Under this conditions the distance between two distinguishable points is approximately 0.16 mm.
- Another report says that the minimal distinguishable distance is 0.47 mm [Kutter, 1999].

Resolution

- Resolution in optics is defined as the inverse value of a maximum viewing angle between the viewer and two proximate points which human cannot distinguish, and so fuse together.
- Human vision has the best resolution for objects which are at a distance of about 25 cm from an eye under illumination of about 500 lux.
- This illumination is provided by a 60 W from a distance 40 cm.
- Under this conditions the distance between two distinguishable points is approximately 0.16 mm.
- Another report says that the minimal distinguishable distance is 0.47 mm [Kutter, 1999].

Resolution

- Resolution in optics is defined as the inverse value of a maximum viewing angle between the viewer and two proximate points which human cannot distinguish, and so fuse together.
- Human vision has the best resolution for objects which are at a distance of about 25 cm from an eye under illumination of about 500 lux.
- This illumination is provided by a 60 W from a distance 40 cm.
- Under this conditions the distance between two distinguishable points is approximately 0.16 mm.
- Another report says that the minimal distinguishable distance is 0.47 mm [Kutter, 1999].

Resolution

- Resolution in optics is defined as the inverse value of a maximum viewing angle between the viewer and two proximate points which human cannot distinguish, and so fuse together.
- Human vision has the best resolution for objects which are at a distance of about 25 cm from an eye under illumination of about 500 lux.
- This illumination is provided by a 60 W from a distance 40 cm.
- Under this conditions the distance between two distinguishable points is approximately 0.16 mm.
- Another report says that the minimal distinguishable distance is 0.47 mm [Kutter, 1999].

Quiz

Given the above two minimal distinguishable distance, what is the resolution in DPI needed for a printer to produce perfect output?
DPI means "Dots Per Inch" ($1\,\mathrm{in} = 2.54\,\mathrm{cm}$).

Visual Illusions

- Human perception of images is prone to many illusions.
- There are many other visual illusions caused by phenomena such as color or motion;
- an Internet search will produce examples easily.

Visual Illusions

- Human perception of images is prone to many illusions.
- There are many other visual illusions caused by phenomena such as color or motion;
- an Internet search will produce examples easily.

Visual Illusions

- Human perception of images is prone to many illusions.
- There are many other visual illusions caused by phenomena such as color or motion;
- an Internet search will produce examples easily.

The Ebbinghaus Circles

- Object borders carry a lot of information.
- Boundaries of objects and simple patterns such as blobs or lines enable adaption effects similar to conditional contrast.
- The Ebbinghaus illusion is a well known example — two circles of the same diameter in the center of images appear to have different sizes.

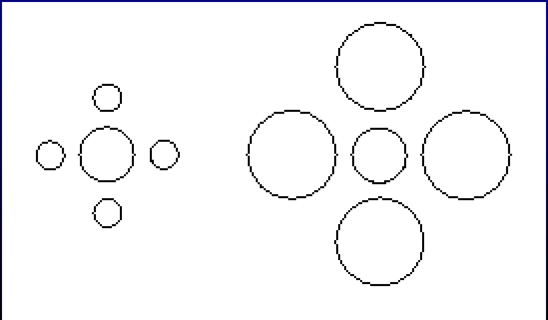

The Ebbinghaus Circles

- Object borders carry a lot of information.
- Boundaries of objects and simple patterns such as blobs or lines enable adaption effects similar to conditional contrast.
- The Ebbinghaus illusion is a well known example — two circles of the same diameter in the center of images appear to have different sizes.

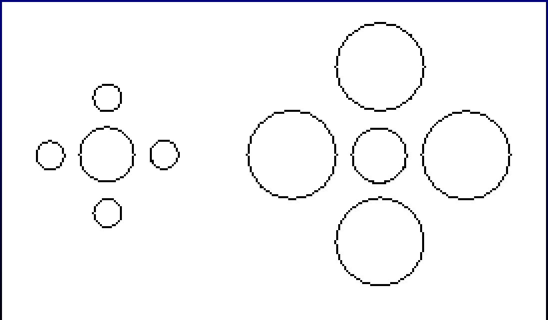

The Ebbinghaus Circles

- Object borders carry a lot of information.
- Boundaries of objects and simple patterns such as blobs or lines enable adaption effects similar to conditional contrast.
- The Ebbinghaus illusion is a well known example — two circles of the same diameter in the center of images appear to have different sizes.

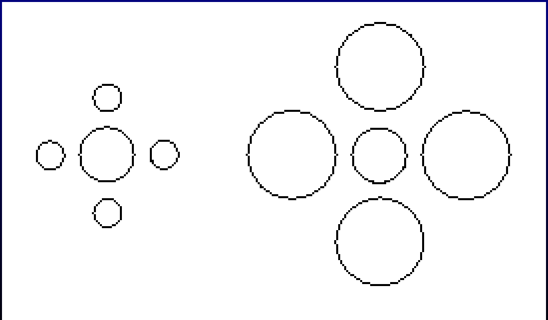

Parallel Lines

- Perception of one dominant shape can be fooled by nearby shapes.
- This figure shows parallel diagonal line segments which are not perceived as parallel.

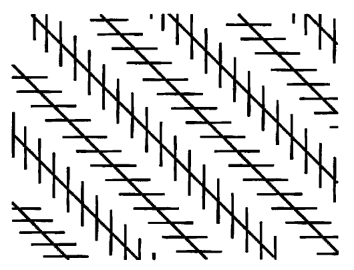

Parallel Lines

- Perception of one dominant shape can be fooled by nearby shapes.
- This figure shows parallel diagonal line segments which are not perceived as parallel.

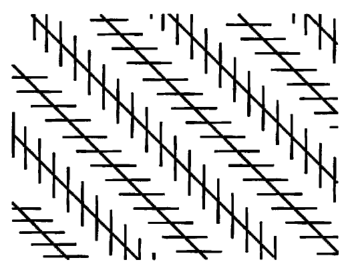

Zigzag Lines

- This figure contains rows of black and white squares which are all parallel.
- However, the vertical zigzag squares disrupt our horizontal perception.

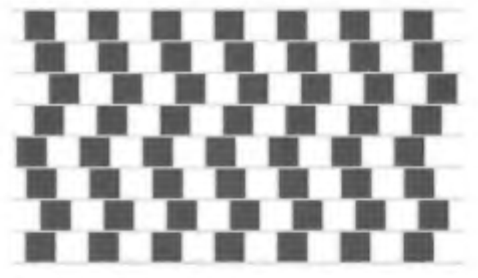

Zigzag Lines

- This figure contains rows of black and white squares which are all parallel.
- However, the vertical zigzag squares disrupt our horizontal perception.

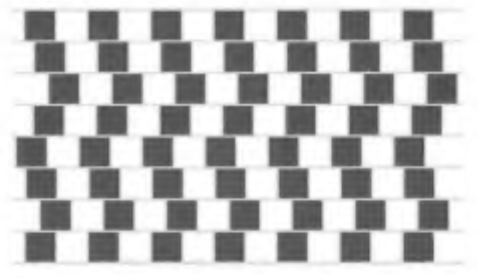

Outline

Image Quality

- An image might be degraded during capture, transmission, or processing.
- Measures of image quality can be used to assess the degree of degradation.
- The quality required naturally depends on the purpose for which an image is used.
- Methods for assessing image quality can be divided into two categories: subjective and objective.

Image Quality

- An image might be degraded during capture, transmission, or processing.
- Measures of image quality can be used to assess the degree of degradation.
- The quality required naturally depends on the purpose for which an image is used.
- Methods for assessing image quality can be divided into two categories: subjective and objective.

Image Quality

- An image might be degraded during capture, transmission, or processing.
- Measures of image quality can be used to assess the degree of degradation.
- The quality required naturally depends on the purpose for which an image is used.
- Methods for assessing image quality can be divided into two categories: subjective and objective.

Image Quality

- An image might be degraded during capture, transmission, or processing.
- Measures of image quality can be used to assess the degree of degradation.
- The quality required naturally depends on the purpose for which an image is used.
- Methods for assessing image quality can be divided into two categories: subjective and objective.

Subjective Quality

- Subjective methods are often used in television technology.
- The ultimate criterion is the perception of a selected group of professional and lay viewers.
- They appraise an image according to a list of criteria and give appropriate marks.

Subjective Quality

- Subjective methods are often used in television technology.
- The ultimate criterion is the perception of a selected group of professional and lay viewers.
- They appraise an image according to a list of criteria and give appropriate marks.

Subjective Quality

- Subjective methods are often used in television technology.
- The ultimate criterion is the perception of a selected group of professional and lay viewers.
- They appraise an image according to a list of criteria and give appropriate marks.

Objective Quality

- Objective quantitative methods measuring image quality are more interesting for our purposes.
- Ideally such a method also provides a good subjective test, and is easy to apply;
- we might then use it as a criterion in parameter optimization.

Objective Quality

- Objective quantitative methods measuring image quality are more interesting for our purposes.
- Ideally such a method also provides a good subjective test, and is easy to apply;
- we might then use it as a criterion in parameter optimization.

Objective Quality

- Objective quantitative methods measuring image quality are more interesting for our purposes.
- Ideally such a method also provides a good subjective test, and is easy to apply;
- we might then use it as a criterion in parameter optimization.

Image Quality: MSE, etc. I

- The quality of the image $f(x, y)$ is usually estimated by comparison with a known reference image $g(x, y)$.
- A synthesized image $g(x, y)$ is often used for this purpose.
- One class of methods uses simple measures such as the mean quadratic difference (or mean squared error, MSE)

$$\mathrm{MSE}(g, f) = \frac{1}{N} \sum_{x,y} (g(x, y) - f(x, y))^2 \qquad (21)$$

where N is the number of pixels.

Image Quality: MSE, etc. II

- The problem here is that it is not possible to distinguish a few big differences from a lot of small differences.
- Instead of the mean quadratic difference, the mean absolute difference or simply maximal absolute difference may be used.
- Correlation between images f and g is another alternative.

Image Quality: SNR

- Signal to noise ratio SNR is also used as a image degradation measure.
- Let $f(x, y)$ be the original image and $f'(x, y)$ be the degraded image, the degree of degradation is measured by

$$\mathrm{SNR}(f', f) = 10 \log_{10} \frac{\sum_{x,y} f(x, y)^2}{\sum_{x,y} (f(x, y) - f'(x, y))^2} \quad \text{(db)} \tag{22}$$

Image Quality: SNR

- Signal to noise ratio SNR is also used as a image degradation measure.
- Let $f(x, y)$ be the original image and $f'(x, y)$ be the degraded image, the degree of degradation is measured by

$$\mathrm{SNR}(f', f) = 10 \log_{10} \frac{\sum_{x,y} f(x,y)^2}{\sum_{x,y} (f(x,y) - f'(x,y))^2} \quad \text{(db)} \tag{22}$$

Image Quality: PSNR

- Peak signal to noise ratio PSNR is another measure in this class.
- PSNR is defined as

$$\mathrm{PSNR}(f', f) = 10 \log_{10} \frac{\max_{x,y} f(x,y)^2}{MSE(f', f)} \quad (23)$$

$$= 10 \log_{10} \frac{N \max_{x,y} f(x,y)^2}{\sum_{x,y}(f(x,y) - f'(x,y))^2} \quad \text{(db)} \quad (24)$$

  where N is the number of pixels.
- Experimentally, a PSNR larger than 32db means invisible visual degradation.

Image Quality: PSNR

- Peak signal to noise ratio PSNR is another measure in this class.
- PSNR is defined as

$$\mathrm{PSNR}(f', f) = 10 \log_{10} \frac{\max_{x,y} f(x,y)^2}{MSE(f', f)} \tag{23}$$

$$= 10 \log_{10} \frac{N \max_{x,y} f(x,y)^2}{\sum_{x,y} (f(x,y) - f'(x,y))^2} \quad \text{(db)} \tag{24}$$

  where N is the number of pixels.
- Experimentally, a PSNR larger than 32db means invisible visual degradation.

Image Quality: PSNR

- Peak signal to noise ratio PSNR is another measure in this class.
- PSNR is defined as

$$\mathrm{PSNR}(f', f) = 10 \log_{10} \frac{\max_{x,y} f(x,y)^2}{MSE(f', f)} \quad (23)$$

$$= 10 \log_{10} \frac{N \max_{x,y} f(x,y)^2}{\sum_{x,y}(f(x,y) - f'(x,y))^2} \quad \text{(db)} \quad (24)$$

where N is the number of pixels.

- Experimentally, a PSNR larger than 32db means invisible visual degradation.

Image Quality Measures

- Measures of image similarity are becoming more important since they may be used in image retrieval from multimedia databases.
- There are many other measures of image similarity based on distance functions [D. R. Wilson and T. R. Martinez, 1997].

Image Quality Measures

- Measures of image similarity are becoming more important since they may be used in image retrieval from multimedia databases.
- There are many other measures of image similarity based on distance functions [D. R. Wilson and T. R. Martinez, 1997].

Outline

Image Noise

- Images are often degraded by random noise.
- Noise can occur during image capture, transmission or processing, and may be dependent on or independent of image content.
- Noise is usually described by its probabilistic characteristics.

Image Noise

- Images are often degraded by random noise.
- Noise can occur during image capture, transmission or processing, and may be dependent on or independent of image content.
- Noise is usually described by its probabilistic characteristics.

Image Noise

- Images are often degraded by random noise.
- Noise can occur during image capture, transmission or processing, and may be dependent on or independent of image content.
- Noise is usually described by its probabilistic characteristics.

White Noise

- White noise — constant power spectrum (its intensity does not decrease with increasing frequency);
- it is frequently applied as a crude approximation of image noise in most cases.
- Its auto-correlation is the delta function. So it is un-correlated at two different instances.
- The advantage is that it simplifies the calculations.

White Noise

- White noise — constant power spectrum (its intensity does not decrease with increasing frequency);
- it is frequently applied as a crude approximation of image noise in most cases.
- Its auto-correlation is the delta function. So it is un-correlated at two different instances.
- The advantage is that it simplifies the calculations.

White Noise

- White noise — constant power spectrum (its intensity does not decrease with increasing frequency);
- it is frequently applied as a crude approximation of image noise in most cases.
- Its auto-correlation is the delta function. So it is un-correlated at two different instances.
- The advantage is that it simplifies the calculations.

White Noise

- White noise — constant power spectrum (its intensity does not decrease with increasing frequency);
- it is frequently applied as a crude approximation of image noise in most cases.
- Its auto-correlation is the delta function. So it is un-correlated at two different instances.
- The advantage is that it simplifies the calculations.

Gaussian Noise

- A special case of noise is Gaussian noise.
- Gaussian noise is a very good approximation of noise that occurs in many practical cases.
- Probability density of the random variable is given by the Gaussian function.
- 1D Gaussian noise — μ is the mean and σ is the standard deviation of the random variable.

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}. \quad (25)$$

Gaussian Noise

- A special case of noise is Gaussian noise.
- Gaussian noise is a very good approximation of noise that occurs in many practical cases.
- Probability density of the random variable is given by the Gaussian function.
- 1D Gaussian noise — μ is the mean and σ is the standard deviation of the random variable.

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}. \quad (25)$$

Gaussian Noise

- A special case of noise is Gaussian noise.
- Gaussian noise is a very good approximation of noise that occurs in many practical cases.
- Probability density of the random variable is given by the Gaussian function.
- 1D Gaussian noise — μ is the mean and σ is the standard deviation of the random variable.

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}. \qquad (25)$$

Gaussian Noise

- A special case of noise is Gaussian noise.
- Gaussian noise is a very good approximation of noise that occurs in many practical cases.
- Probability density of the random variable is given by the Gaussian function.
- 1D Gaussian noise — μ is the mean and σ is the standard deviation of the random variable.

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \mathbf{e}^{-\frac{(x-\mu)^2}{2\sigma^2}}. \tag{25}$$

Noise Type: additive

- Noise may be

  additive: the noise ν and image signal g are independent

$$f(x, y) = g(x, y) + \nu(x, y). \tag{26}$$

- During image transmission, noise is usually independent of the image signal occurs.
- The degradation can be modeled as additive noise.

Noise Type: additive

- Noise may be

  additive: the noise ν and image signal g are independent

$$f(x, y) = g(x, y) + \nu(x, y). \qquad (26)$$

- During image transmission, noise is usually independent of the image signal occurs.
- The degradation can be modeled as additive noise.

Noise Type: additive

- Noise may be

  additive: the noise ν and image signal g are independent

$$f(x, y) = g(x, y) + \nu(x, y). \tag{26}$$

- During image transmission, noise is usually independent of the image signal occurs.
- The degradation can be modeled as additive noise.

Noise Type: additive

- Noise may be

  additive: the noise ν and image signal g are independent

$$f(x,y) = g(x,y) + \nu(x,y). \tag{26}$$

- During image transmission, noise is usually independent of the image signal occurs.
- The degradation can be modeled as additive noise.

Noise Type: multiplicative and impulse

- Noise may be

multiplicative
: the noise is a function of signal magnitude

$$f(x,y) = g(x,y) + \nu(x,y)g(x,y) \quad (27)$$
$$= g(x,y)(1+\nu(x,y)) \quad (28)$$
$$= g(x,y)n(x,y). \quad (29)$$

impulse
: an image is corrupted with individual noisy pixels whose brightness differs significantly from that of the neighborhood.

Noise Type: multiplicative and impulse

- Noise may be

  multiplicative — the noise is a function of signal magnitude

$$f(x, y) = g(x, y) + \nu(x, y)g(x, y) \quad (27)$$
$$= g(x, y)(1 + \nu(x, y)) \quad (28)$$
$$= g(x, y)n(x, y). \quad (29)$$

  impulse — an image is corrupted with individual noisy pixels whose brightness differs significantly from that of the neighborhood.

Noise Type: multiplicative and impulse

- Noise may be

  multiplicative the noise is a function of signal magnitude

$$
\begin{aligned}
f(x, y) &= g(x, y) + \nu(x, y)g(x, y) && (27)\\
&= g(x, y)(1 + \nu(x, y)) && (28)\\
&= g(x, y)n(x, y). && (29)
\end{aligned}
$$

  impulse an image is corrupted with individual noisy pixels whose brightness differs significantly from that of the neighborhood.

Noise: discussions

- The term “salt and pepper noise” is used to describe saturated impulsive noise — an image corrupted with white and/or black pixel is an example.
- The problem of suppressing noise in images is addressed in subsequent lectures of this course.

Noise: discussions

- The term “salt and pepper noise” is used to describe saturated impulsive noise — an image corrupted with white and/or black pixel is an example.
- The problem of suppressing noise in images is addressed in subsequent lectures of this course.

Simulation of Gaussian Noise

- We first consider the simulation of Gaussian noise.
- Theorem

(Box-Muller Method) Let U_1 and U_2 be i.i.d (independent identical distributed) uniformly distributed random variables on $(0, 1)$. Then the random variables

$$N_1 = \sqrt{-2 \log U_1} \cos(2\pi U_2) \quad (30)$$

$$N_2 = \sqrt{-2 \log U_1} \sin(2\pi U_2) \quad (31)$$

are independent standard Gaussian.

Simulation of Gaussian Noise

- We first consider the simulation of Gaussian noise.

- **Theorem**

  (Box-Muller Method) Let U_1 and U_2 be i.i.d (independent identical distributed) uniformly distributed random variables on $(0, 1)$. Then the random variables

$$N_1 = \sqrt{-2 \log U_1} \cos(2\pi U_2) \tag{30}$$

$$N_2 = \sqrt{-2 \log U_1} \sin(2\pi U_2) \tag{31}$$

are independent standard Gaussian.

Simulation of Gaussian Noise

- We first consider the simulation of Gaussian noise.

- Theorem

(Box-Muller Method) Let U_1 and U_2 be i.i.d (independent identical distributed) uniformly distributed random variables on $(0, 1)$. Then the random variables

$$N_1 = \sqrt{-2 \log U_1} \cos(2\pi U_2) \tag{30}$$
$$N_2 = \sqrt{-2 \log U_1} \sin(2\pi U_2) \tag{31}$$

are independent standard Gaussian.

Proof of Box-Muller Method I

- To prove it we need the following

- **Theorem**
  (Transformation Theorem for Densities) Let Z_1, Z_2 and U_1, U_2 be random variables. Assume
  - (U_1, U_2) *takes values in the open set* G' *of* $\mathbf{R}^2$ *and has density* f *on* G'*;*
  - (Z_1, Z_2) *takes values in the open set* G *of* $\mathbf{R}^2$*;*
  - $\varphi : G \longmapsto G'$ *is a continuously differentiable bijection with continuously differentiable inverse* $\varphi^{-1} : G' = \varphi(G) \longmapsto G$.

  Given
  $$\begin{pmatrix} U_1 \\ U_2 \end{pmatrix} = \varphi \begin{pmatrix} Z_1 \\ Z_2 \end{pmatrix}. \tag{32}$$
  the random vector (Z_1, Z_2) *on* G *has the density*
  $$g(z) = f \circ \varphi(z) |J_\varphi(z)| \tag{33}$$
  where $J_\varphi(z) = \frac{\partial(\varphi_1, \varphi_2)}{\partial(z_1, z_2)}$ *is the Jacobian of* φ.

Proof of Box-Muller Method I

- To prove it we need the following

- **Theorem**

(Transformation Theorem for Densities) Let Z_1, Z_2 and U_1, U_2 be random variables. Assume

  - *(U_1, U_2) takes values in the open set G' of $\mathbf{R}^2$ and has density f on G';*
  - *(Z_1, Z_2) takes values in the open set G of $\mathbf{R}^2$;*
  - *$\varphi : G \longmapsto G'$ is a continuously differentiable bijection with continuously differentiable inverse $\varphi^{-1} : G' = \varphi(G) \longmapsto G$.*

Given

$$\begin{pmatrix} U_1 \\ U_2 \end{pmatrix} = \varphi \begin{pmatrix} Z_1 \\ Z_2 \end{pmatrix}. \tag{32}$$

the random vector (Z_1, Z_2) on G has the density

$$g(z) = f \circ \varphi(z) |J_\varphi(z)| \tag{33}$$

where $J_\varphi(z) = \frac{\partial(\varphi_1, \varphi_2)}{\partial(z_1, z_2)}$ is the Jacobian of φ.

Proof of Box-Muller Method I

- To prove it we need the following

- Theorem

(Transformation Theorem for Densities) Let Z_1, Z_2 and U_1, U_2 be random variables. Assume

  - *(U_1, U_2) takes values in the open set G' of $\mathbf{R}^2$ and has density f on G';*
  - *(Z_1, Z_2) takes values in the open set G of $\mathbf{R}^2$;*
  - *$\varphi : G \longmapsto G'$ is a continuously differentiable bijection with continuously differentiable inverse $\varphi^{-1} : G' = \varphi(G) \longmapsto G$.*

Given

$$\begin{pmatrix} U_1 \\ U_2 \end{pmatrix} = \varphi \begin{pmatrix} Z_1 \\ Z_2 \end{pmatrix}. \tag{32}$$

the random vector (Z_1, Z_2) on G has the density

$$g(z) = f \circ \varphi(z) |J_\varphi(z)| \tag{33}$$

where $J_\varphi(z) = \frac{\partial(\varphi_1, \varphi_2)}{\partial(z_1, z_2)}$ is the Jacobian of φ.

Proof of Box-Muller Method I

- To prove it we need the following

- Theorem

(Transformation Theorem for Densities) Let Z_1, Z_2 and U_1, U_2 be random variables. Assume

  - (U_1, U_2) *takes values in the open set* G' *of* $\mathbf{R}^2$ *and has density* f *on* G';
  - (Z_1, Z_2) *takes values in the open set* G *of* $\mathbf{R}^2$;
  - $\varphi : G \longmapsto G'$ *is a continuously differentiable bijection with continuously differentiable inverse* $\varphi^{-1} : G' = \varphi(G) \longmapsto G$.

Given

$$\begin{pmatrix} U_1 \\ U_2 \end{pmatrix} = \varphi \begin{pmatrix} Z_1 \\ Z_2 \end{pmatrix}. \tag{32}$$

the random vector (Z_1, Z_2) *on* G *has the density*

$$g(z) = f \circ \varphi(z) |J_\varphi(z)| \tag{33}$$

where $J_\varphi(z) = \frac{\partial(\varphi_1, \varphi_2)}{\partial(z_1, z_2)}$ *is the Jacobian of* φ.

Proof of Box-Muller Method I

- To prove it we need the following

- Theorem
  (Transformation Theorem for Densities) Let Z_1, Z_2 and U_1, U_2 be random variables. Assume
  - (U_1, U_2) *takes values in the open set* G' *of* $\mathbf{R}^2$ *and has density* f *on* G'*;*
  - (Z_1, Z_2) *takes values in the open set* G *of* $\mathbf{R}^2$*;*
  - $\varphi : G \longmapsto G'$ *is a continuously differentiable bijection with continuously differentiable inverse* $\varphi^{-1} : G' = \varphi(G) \longmapsto G$.

Given

$$\begin{pmatrix} U_1 \\ U_2 \end{pmatrix} = \varphi \begin{pmatrix} Z_1 \\ Z_2 \end{pmatrix}. \tag{32}$$

the random vector (Z_1, Z_2) *on* G *has the density*

$$g(z) = f \circ \varphi(z) |J_\varphi(z)| \tag{33}$$

where $J_\varphi(z) = \frac{\partial(\varphi_1, \varphi_2)}{\partial(z_1, z_2)}$ *is the Jacobian of* φ.

Proof of Box-Muller Method I

- To prove it we need the following

- Theorem

(Transformation Theorem for Densities) Let Z_1, Z_2 and U_1, U_2 be random variables. Assume

- (U_1, U_2) *takes values in the open set G' of $\mathbf{R}^2$ and has density f on G';*
- (Z_1, Z_2) *takes values in the open set G of $\mathbf{R}^2$;*
- $\varphi : G \longmapsto G'$ *is a continuously differentiable bijection with continuously differentiable inverse* $\varphi^{-1} : G' = \varphi(G) \longmapsto G$.

Given

$$\begin{pmatrix} U_1 \\ U_2 \end{pmatrix} = \varphi \begin{pmatrix} Z_1 \\ Z_2 \end{pmatrix}. \tag{32}$$

the random vector (Z_1, Z_2) on G has the density

$$g(z) = f \circ \varphi(z) |J_\varphi(z)| \tag{33}$$

where $J_\varphi(z) = \frac{\partial(\varphi_1, \varphi_2)}{\partial(z_1, z_2)}$ is the Jacobian of φ.

Proof of Box-Muller Method II

- Frst we determine the map φ from the last theorem.
- We have

$$N_1^2 = -2\log U_1 \cos^2(2\pi U_2) \tag{34}$$
$$N_2^2 = -2\log U_1 \sin^2(2\pi U_2). \tag{35}$$

- Hence

$$N_1^2 + N_2^2 = -2\log U_1 \tag{36}$$

and

$$U_1 = \mathbf{e}^{-\frac{N_1^2+N_2^2}{2}}. \tag{37}$$

- Moreover, by (30),

$$\frac{N_2}{N_1} = \tan(2\pi U_2), \tag{38}$$

i.e.,

$$U_2 = \frac{1}{2\pi}\arctan\left(\frac{N_2}{N_1}\right). \tag{39}$$

Proof of Box-Muller Method II

- Frst we determine the map φ from the last theorem.
- We have

$$N_1^2 = -2\log U_1 \cos^2(2\pi U_2) \tag{34}$$

$$N_2^2 = -2\log U_1 \sin^2(2\pi U_2). \tag{35}$$

- Hence

$$N_1^2 + N_2^2 = -2\log U_1 \tag{36}$$

and

$$U_1 = \mathbf{e}^{-\frac{N_1^2+N_2^2}{2}}. \tag{37}$$

- Moreover, by (30),

$$\frac{N_2}{N_1} = \tan(2\pi U_2), \tag{38}$$

i.e.,

$$U_2 = \frac{1}{2\pi}\arctan\left(\frac{N_2}{N_1}\right). \tag{39}$$

Proof of Box-Muller Method II

- Frst we determine the map φ from the last theorem.
- We have

$$N_1^2 = -2\log U_1 \cos^2(2\pi U_2) \quad (34)$$

$$N_2^2 = -2\log U_1 \sin^2(2\pi U_2). \quad (35)$$

- Hence

$$N_1^2 + N_2^2 = -2\log U_1 \quad (36)$$

and

$$U_1 = \mathbf{e}^{-\frac{N_1^2+N_2^2}{2}}. \quad (37)$$

- Moreover, by (30),

$$\frac{N_2}{N_1} = \tan(2\pi U_2), \quad (38)$$

i.e.,

$$U_2 = \frac{1}{2\pi}\arctan\left(\frac{N_2}{N_1}\right). \quad (39)$$

Proof of Box-Muller Method II

- Frst we determine the map φ from the last theorem.
- We have

$$N_1^2 = -2\log U_1 \cos^2(2\pi U_2) \tag{34}$$
$$N_2^2 = -2\log U_1 \sin^2(2\pi U_2). \tag{35}$$

- Hence

$$N_1^2 + N_2^2 = -2\log U_1 \tag{36}$$

and

$$U_1 = \mathbf{e}^{-\frac{N_1^2+N_2^2}{2}}. \tag{37}$$

- Moreover, by (30),

$$\frac{N_2}{N_1} = \tan(2\pi U_2), \tag{38}$$

i.e.,

$$U_2 = \frac{1}{2\pi}\arctan\left(\frac{N_2}{N_1}\right). \tag{39}$$

Proof of Box-Muller Method III

- Hence

$$\varphi(z_1, z_2) = \begin{pmatrix} e^{-\frac{z_1^2+z_2^2}{2}} \\ \frac{1}{2\pi} \arctan(\frac{z_2}{z_1}) \end{pmatrix} \tag{40}$$

- The transform domains are

$$G = \mathbf{R}^2 \setminus [\{z_1 = 0\} \cup \{z_2 = 0 \text{ and } z_1 > 0\}], \tag{41}$$

$$G' = (0, 1) \times (0, 1) \setminus \{u_2 = \frac{1}{4} \text{ or } u_2 = \frac{3}{4}\}. \tag{42}$$

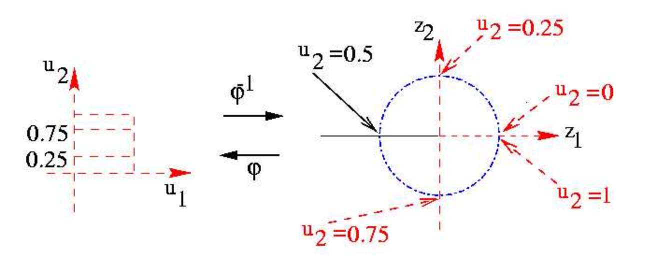

Proof of Box-Muller Method III

- Hence

$$\varphi(z_1, z_2) = \begin{pmatrix} \mathbf{e}^{-\frac{z_1^2+z_2^2}{2}} \\ \frac{1}{2\pi} \arctan(\frac{z_2}{z_1}) \end{pmatrix} \quad (40)$$

- The transform domains are

$$G = \mathbf{R}^2 \setminus [\{z_1 = 0\} \cup \{z_2 = 0 \text{ and } z_1 > 0\}], \quad (41)$$

$$G' = (0,1) \times (0,1) \setminus \{u_2 = \frac{1}{4} \text{ or } u_2 = \frac{3}{4}\}. \quad (42)$$

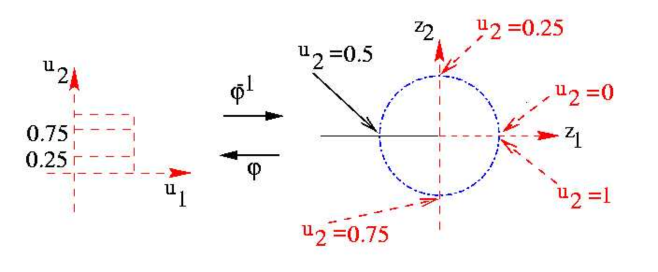

Proof of Box-Muller Method IV

- The partial derivatives of φ are

$$\frac{\partial \varphi_1}{\partial z_1}(z) = -z_1 \mathbf{e}^{-\frac{z_1^2+z_2^2}{2}}, \qquad \frac{\partial \varphi_1}{\partial z_2}(z) = -z_2 \mathbf{e}^{-\frac{z_1^2+z_2^2}{2}} \tag{43}$$

$$\frac{\partial \varphi_1}{\partial z_1}(z) = \frac{1}{2\pi}\frac{-z_2}{z_1^2+z_2^2}, \qquad \frac{\partial \varphi_1}{\partial z_2}(z) = \frac{1}{2\pi}\frac{z_1}{z_1^2+z_2^2}. \tag{44}$$

- It follows that

$$J_\varphi(z) = \frac{1}{2\pi}\mathbf{e}^{-\frac{z_1^2+z_2^2}{2}} = \frac{1}{\sqrt{2\pi}}\mathbf{e}^{-\frac{z_1^2}{2}} \cdot \frac{1}{\sqrt{2\pi}}\mathbf{e}^{-\frac{z_2^2}{2}}. \tag{45}$$

Proof of Box-Muller Method IV

- The partial derivatives of φ are

$$\frac{\partial\varphi_1}{\partial z_1}(z) = -z_1 \mathbf{e}^{-\frac{z_1^2+z_2^2}{2}}, \qquad \frac{\partial\varphi_1}{\partial z_2}(z) = -z_2 \mathbf{e}^{-\frac{z_1^2+z_2^2}{2}} \tag{43}$$

$$\frac{\partial\varphi_1}{\partial z_1}(z) = \frac{1}{2\pi}\frac{-z_2}{z_1^2+z_2^2}, \qquad \frac{\partial\varphi_1}{\partial z_2}(z) = \frac{1}{2\pi}\frac{z_1}{z_1^2+z_2^2}. \tag{44}$$

- It follows that

$$J_\varphi(z) = \frac{1}{2\pi}\mathbf{e}^{-\frac{z_1^2+z_2^2}{2}} = \frac{1}{\sqrt{2\pi}}\mathbf{e}^{-\frac{z_1^2}{2}} \cdot \frac{1}{\sqrt{2\pi}}\mathbf{e}^{-\frac{z_2^2}{2}}. \tag{45}$$

Proof of Box-Muller Method V

- Since (U_1, U_2) has density $\chi_{(0,1)\times(0,1)}$, which is identically 1 on $\chi_{(0,1)\times(0,1)}$, (N_1, N_2) has density

$$\frac{1}{\sqrt{2\pi}}e^{-\frac{z_1^2}{2}} \cdot \frac{1}{\sqrt{2\pi}}e^{-\frac{z_2^2}{2}} \tag{46}$$

on G. Therefore they are independent Gaussian variables. $\square$

The Rejection Method I

- This method uses an auxiliary density for generation of random quantities from distributions not amenable to analytic treatment [Press et al., 1992].
- Consider the generation of samples from a density π.
- Consider also an auxiliary density q for which we know how to generate samples.
- The idea is to use q to make samples from π.
- The method is general enough to generate samples from π without knowing the complete expression for π.

The Rejection Method I

- This method uses an auxiliary density for generation of random quantities from distributions not amenable to analytic treatment [Press et al., 1992].
- Consider the generation of samples from a density π.
- Consider also an auxiliary density q for which we know how to generate samples.
- The idea is to use q to make samples from π.
- The method is general enough to generate samples from π without knowing the complete expression for π.

The Rejection Method I

- This method uses an auxiliary density for generation of random quantities from distributions not amenable to analytic treatment [Press et al., 1992].
- Consider the generation of samples from a density π.
- Consider also an auxiliary density q for which we know how to generate samples.
- The idea is to use q to make samples from π.
- The method is general enough to generate samples from π without knowing the complete expression for π.

The Rejection Method I

- This method uses an auxiliary density for generation of random quantities from distributions not amenable to analytic treatment [Press et al., 1992].
- Consider the generation of samples from a density π.
- Consider also an auxiliary density q for which we know how to generate samples.
- The idea is to use q to make samples from π.
- The method is general enough to generate samples from π without knowing the complete expression for π.

The Rejection Method I

- This method uses an auxiliary density for generation of random quantities from distributions not amenable to analytic treatment [Press et al., 1992].
- Consider the generation of samples from a density π.
- Consider also an auxiliary density q for which we know how to generate samples.
- The idea is to use q to make samples from π.
- The method is general enough to generate samples from π without knowing the complete expression for π.

The Rejection Method II

- It is extremely common in statistics to encounter such situations where the kernel of π is known but the constant ensuring it integrals to 1 cannot be obtained analytically.
- The only mathematical restriction over q is that there must exist a constant A such that

$$\pi(x) \leq Aq(x) \tag{47}$$

for every possible value of x.

The Rejection Method II

- It is extremely common in statistics to encounter such situations where the kernel of π is known but the constant ensuring it integrals to 1 cannot be obtained analytically.
- The only mathematical restriction over q is that there must exist a constant A such that

$$\pi(x) \leq Aq(x) \tag{47}$$

for every possible value of x.

The Rejection Method III

- The method consists of independently drawing X from q and $U \sim U[0,1]$ and accepting X as a sample generated from π if

$$AUq(X) \leq \pi(X). \tag{48}$$

Otherwise X is not accepted as a sample from π and the process must be reinitialized until a value X is accepted.

- Hence the name of the method, which is also known as the acceptance/rejection method.

The Rejection Method III

- The method consists of independently drawing X from q and $U \sim U[0,1]$ and accepting X as a sample generated from π if

$$AUq(X) \leq \pi(X). \tag{48}$$

Otherwise X is not accepted as a sample from π and the process must be reinitialized until a value X is accepted.

- Hence the name of the method, which is also known as the acceptance/rejection method.

Proof of the Rejection Method I

- To prove the rejection procedure effectively generates samples from π, we need to show that the density function of X is a constant multiple of π, when conditioned by the acceptance condition (48).
- The joint density of (X, U) is $f(x, u) = q(x)\chi_{[0,1]}(u)$, by the independence between X and U and the uniformity of U.
- The conditional probability is computed by the following formula:

$$\mathbf{Pr}(A|B) = \frac{\mathbf{Pr}(A, B)}{\mathbf{Pr}(B)}. \quad (49)$$

Proof of the Rejection Method I

- To prove the rejection procedure effectively generates samples from π, we need to show that the density function of X is a constant multiple of π, when conditioned by the acceptance condition (48).
- The joint density of (X, U) is $f(x, u) = q(x)\chi_{[0,1]}(u)$, by the independence between X and U and the uniformity of U.
- The conditional probability is computed by the following formula:

$$\mathbf{Pr}(A|B) = \frac{\mathbf{Pr}(A, B)}{\mathbf{Pr}(B)}. \tag{49}$$

Proof of the Rejection Method I

- To prove the rejection procedure effectively generates samples from π, we need to show that the density function of X is a constant multiple of π, when conditioned by the acceptance condition (48).
- The joint density of (X, U) is $f(x, u) = q(x)\chi_{[0,1]}(u)$, by the independence between X and U and the uniformity of U.
- The conditional probability is computed by the following formula:

$$\mathbf{Pr}(A|B) = \frac{\mathbf{Pr}(A, B)}{\mathbf{Pr}(B)}. \tag{49}$$

Proof of the Rejection Method II

$$\mathbf{Pr}(X \le t | AUq(X) \le \pi(X)) = \mathbf{Pr}(X \le t | U \le \frac{\pi(X)}{Aq(X)})$$

$$= \frac{\mathbf{Pr}(X \le t, U \le \frac{\pi(X)}{Aq(X)})}{\mathbf{Pr}(U \le \frac{\pi(X)}{Aq(X)})}$$

$$= \frac{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} q(x)\chi_{[0,1]}(u)\chi_{\left\{x \le t, u \le \frac{\pi(x)}{Aq(x)}\right\}}(x,u)\,dxdu}{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} q(x)\chi_{[0,1]}(u)\chi_{\left\{u \le \frac{\pi(x)}{Aq(x)}\right\}}(x,u)\,dxdu}$$

$$= \frac{\int_{-\infty}^{t} dx \int_0^{\frac{\pi(x)}{Aq(x)}} q(x)\chi_{[0,1]}(u)\,du}{\int_{-\infty}^{\infty} dx \int_0^{\frac{\pi(x)}{Aq(x)}} q(x)\chi_{[0,1]}(u)\,du}$$

$$= \frac{\int_{-\infty}^{t} \frac{\pi(x)}{Aq(x)} q(x)\,dx}{\int_{-\infty}^{\infty} \frac{\pi(x)}{Aq(x)} q(x)\,dx} = \frac{\int_{-\infty}^{t} \pi(x)\,dx}{\int_{-\infty}^{\infty} \pi(x)\,dx}.$$

Extension of the Rejection Method

- The samples to be generated can have any form: scalar, vector or matrix.
- In each case, the rejection step is based on a comparison of densities with aid of a scalar uniform random variable.

Extension of the Rejection Method

- The samples to be generated can have any form: scalar, vector or matrix.
- In each case, the rejection step is based on a comparison of densities with aid of a scalar uniform random variable.

Discussions on the Rejection Method

- q should be a density that is easy to draw samples from.
- The overall acceptance probability is

$$\mathbf{Pr}\left(U \leq \frac{\pi(X)}{Aq(X)}\right) = \frac{1}{A}\int_{-\infty}^{\infty} \pi(x)\, dx \qquad (50)$$

- Hence, A must be chosen as close as possible to $\int_{-\infty}^{\infty} \pi(x)\, dx$.

Discussions on the Rejection Method

- q should be a density that is easy to draw samples from.
- The overall acceptance probability is

$$\mathbf{Pr}\left(U \leq \frac{\pi(X)}{Aq(X)}\right) = \frac{1}{A} \int_{-\infty}^{\infty} \pi(x)\, dx \tag{50}$$

- Hence, A must be chosen as close as possible to $\int_{-\infty}^{\infty} \pi(x)\, dx$.

Discussions on the Rejection Method

- q should be a density that is easy to draw samples from.
- The overall acceptance probability is

$$\mathbf{Pr}\left(U \leq \frac{\pi(X)}{Aq(X)}\right) = \frac{1}{A}\int_{-\infty}^{\infty} \pi(x)\,dx \qquad (50)$$

- Hence, A must be chosen as close as possible to $\int_{-\infty}^{\infty} \pi(x)\,dx$.

Rejection Method for Truncated Distributions

- A special case of rejection sampling is given for truncated distributions.
- Let q be any density and π be its truncation to the region C, i.e., $\pi = q\chi_C \leq q$.
- Taking $A = 1$, the acceptance condition is

$$Uq(X) \leq \pi(X) = q\chi_C \qquad (51)$$

  which is satisfied, almost surely, if and only if $X \in C$.
- Hence, to generate sample from q restricted to C, one simply has to draw sample from q and accept it if and only if it is in C.

Rejection Method for Truncated Distributions

- A special case of rejection sampling is given for truncated distributions.
- Let q be any density and π be its truncation to the region C, i.e., $\pi = q\chi_C \leq q$.
- Taking $A = 1$, the acceptance condition is

$$Uq(X) \leq \pi(X) = q\chi_C \qquad (51)$$

which is satisfied, almost surely, if and only if $X \in C$.
- Hence, to generate sample from q restricted to C, one simply has to draw sample from q and accept it if and only if it is in C.

Rejection Method for Truncated Distributions

- A special case of rejection sampling is given for truncated distributions.
- Let q be any density and π be its truncation to the region C, i.e., $\pi = q\chi_C \leq q$.
- Taking $A = 1$, the acceptance condition is

$$Uq(X) \leq \pi(X) = q\chi_C \tag{51}$$

which is satisfied, almost surely, if and only if $X \in C$.
- Hence, to generate sample from q restricted to C, one simply has to draw sample from q and accept it if and only if it is in C.

Rejection Method for Truncated Distributions

- A special case of rejection sampling is given for truncated distributions.
- Let q be any density and π be its truncation to the region C, i.e., $\pi = q\chi_C \le q$.
- Taking $A = 1$, the acceptance condition is

$$Uq(X) \le \pi(X) = q\chi_C \tag{51}$$

which is satisfied, almost surely, if and only if $X \in C$.

- Hence, to generate sample from q restricted to C, one simply has to draw sample from q and accept it if and only if it is in C.

Polar Method

- This a variant of the Box-Muller method to generate standard normal deviates. It is due to G. Marsaglia.
- It is substantially faster than the Box-Muller method since it avoids the calculation of the trigonometric functions (but still slower than other methods, [Knuth, 1981]) and it has essential perfect accuracy.

Polar Method

- This a variant of the Box-Muller method to generate standard normal deviates. It is due to G. Marsaglia.
- It is substantially faster than the Box-Muller method since it avoids the calculation of the trigonometric functions (but still slower than other methods, [Knuth, 1981]) and it has essential perfect accuracy.

Box-Muller Method Revisited

- The Box-Muller method may be rephrased as follows:
  - given (W, Θ) uniformly distributed on $[0,1] \times [0, 2\pi]$;
  - the variables

$$N_1 = \sqrt{-2 \log W} \cos(\Theta) \tag{52}$$

$$N_2 = \sqrt{-2 \log W} \sin(\Theta) \tag{53}$$

  are independent standard Gaussian.
- The rejection method allows us to sample directly from $\sqrt{W} \cos \Theta$ and $\sqrt{W} \sin \Theta$, thus avoiding to calculate the sines and cosines.

Box-Muller Method Revisited

- The Box-Muller method may be rephrased as follows:
  - given (W, Θ) uniformly distributed on $[0, 1] \times [0, 2\pi]$;
  - the variables

$$N_1 = \sqrt{-2 \log W} \cos(\Theta) \tag{52}$$

$$N_2 = \sqrt{-2 \log W} \sin(\Theta) \tag{53}$$

    are independent standard Gaussian.
- The rejection method allows us to sample directly from $\sqrt{W} \cos \Theta$ and $\sqrt{W} \sin \Theta$, thus avoiding to calculate the sines and cosines.

Box-Muller Method Revisited

- The Box-Muller method may be rephrased as follows:
  - given (W, Θ) uniformly distributed on $[0, 1] \times [0, 2\pi]$;
  - the variables

$$N_1 = \sqrt{-2 \log W} \cos(\Theta) \tag{52}$$

$$N_2 = \sqrt{-2 \log W} \sin(\Theta) \tag{53}$$

  are independent standard Gaussian.
- The rejection method allows us to sample directly from $\sqrt{W} \cos \Theta$ and $\sqrt{W} \sin \Theta$, thus avoiding to calculate the sines and cosines.

Box-Muller Method Revisited

- The Box-Muller method may be rephrased as follows:
  - given (W, Θ) uniformly distributed on $[0, 1] \times [0, 2\pi]$;
  - the variables

$$N_1 = \sqrt{-2 \log W} \cos(\Theta) \quad (52)$$

$$N_2 = \sqrt{-2 \log W} \sin(\Theta) \quad (53)$$

  are independent standard Gaussian.
- The rejection method allows us to sample directly from $\sqrt{W} \cos \Theta$ and $\sqrt{W} \sin \Theta$, thus avoiding to calculate the sines and cosines.

Sampling $\sqrt{W}\cos\Theta$ and $\sqrt{W}\sin\Theta$

- Given (Z_1, Z_2) uniformly distributed on the unit disk;
- $Z_1 = R\cos\Theta$ and $Z_2 = R\sin\Theta$ in polar coordinates R, Θ;
- Then $W = R^2$ and Θ have joint density

$$\frac{1}{\pi}\chi_{\{|r|\leq 1\}}(\sqrt{W}\cos\Theta, \sqrt{W}\sin\Theta)\left|\frac{\partial Z_1, Z_2}{\partial W, \Theta}\right| \quad (54)$$

$$= \frac{1}{\pi}\chi_{(0,1]\times[0,2\pi)}\left|\begin{pmatrix} \frac{\cos\Theta}{2\sqrt{W}} & \frac{\sin\Theta}{2\sqrt{W}} \\ -\sqrt{W}\sin\Theta & \sqrt{W}\cos\Theta \end{pmatrix}\right| \quad (55)$$

$$= \frac{1}{2\pi}\chi_{[0,1]\times[0,2\pi)} \quad (56)$$

  on $(0, 1] \times [0, 2\pi)$;
- Hence (W, Θ) constructed above is uniform and independent on $(0, 1] \times [0, 2\pi)$;

Sampling $\sqrt{W}\cos\Theta$ and $\sqrt{W}\sin\Theta$

- Given (Z_1, Z_2) uniformly distributed on the unit disk;
- $Z_1 = R\cos\Theta$ and $Z_2 = R\sin\Theta$ in polar coordinates R, Θ;
- Then $W = R^2$ and Θ have joint density

$$\frac{1}{\pi}\chi_{\{|r|\le 1\}}(\sqrt{W}\cos\Theta, \sqrt{W}\sin\Theta)\left|\frac{\partial Z_1, Z_2}{\partial W, \Theta}\right| \tag{54}$$

$$=\frac{1}{\pi}\chi_{(0,1]\times[0,2\pi)}\left|\begin{pmatrix} \frac{\cos\Theta}{2\sqrt{W}} & \frac{\sin\Theta}{2\sqrt{W}} \\ -\sqrt{W}\sin\Theta & \sqrt{W}\cos\Theta \end{pmatrix}\right| \tag{55}$$

$$=\frac{1}{2\pi}\chi_{[0,1]\times[0,2\pi)} \tag{56}$$

  on $(0,1]\times[0,2\pi)$;
- Hence (W, Θ) constructed above is uniform and independent on $(0,1]\times[0,2\pi)$;

Sampling $\sqrt{W}\cos\Theta$ and $\sqrt{W}\sin\Theta$

- Given (Z_1, Z_2) uniformly distributed on the unit disk;
- $Z_1 = R\cos\Theta$ and $Z_2 = R\sin\Theta$ in polar coordinates R, Θ;
- Then $W = R^2$ and Θ have joint density

$$\frac{1}{\pi}\chi_{\{|r|\leq 1\}}(\sqrt{W}\cos\Theta, \sqrt{W}\sin\Theta)\left|\frac{\partial Z_1, Z_2}{\partial W, \Theta}\right| \tag{54}$$

$$=\frac{1}{\pi}\chi_{(0,1]\times[0,2\pi)}\left|\begin{pmatrix} \frac{\cos\Theta}{2\sqrt{W}} & \frac{\sin\Theta}{2\sqrt{W}} \\ -\sqrt{W}\sin\Theta & \sqrt{W}\cos\Theta \end{pmatrix}\right| \tag{55}$$

$$=\frac{1}{2\pi}\chi_{[0,1]\times[0,2\pi)} \tag{56}$$

  on $(0, 1] \times [0, 2\pi)$;
- Hence (W, Θ) constructed above is uniform and independent on $(0, 1] \times [0, 2\pi)$;

Sampling $\sqrt{W}\cos\Theta$ and $\sqrt{W}\sin\Theta$

- Given (Z_1, Z_2) uniformly distributed on the unit disk;
- $Z_1 = R\cos\Theta$ and $Z_2 = R\sin\Theta$ in polar coordinates R, Θ;
- Then $W = R^2$ and Θ have joint density

$$\frac{1}{\pi}\chi_{\{|r|\le 1\}}(\sqrt{W}\cos\Theta, \sqrt{W}\sin\Theta)\left|\frac{\partial Z_1, Z_2}{\partial W, \Theta}\right| \tag{54}$$

$$= \frac{1}{\pi}\chi_{(0,1]\times[0,2\pi)}\left|\begin{pmatrix} \frac{\cos\Theta}{2\sqrt{W}} & \frac{\sin\Theta}{2\sqrt{W}} \\ -\sqrt{W}\sin\Theta & \sqrt{W}\cos\Theta \end{pmatrix}\right| \tag{55}$$

$$= \frac{1}{2\pi}\chi_{[0,1]\times[0,2\pi)} \tag{56}$$

  on $(0, 1] \times [0, 2\pi)$;
- Hence (W, Θ) constructed above is uniform and independent on $(0, 1] \times [0, 2\pi)$;

Box-Muller Method in Polar Representation

- Clearly $W = Z_1^2 + Z_2^2$ and $\cos\Theta = \frac{Z_1}{\sqrt{W}}$, and $\sin\Theta = \frac{Z_2}{\sqrt{W}}$;
- Then the Box-Muller transform can be written as

$$N_1 = \sqrt{\frac{-2\log W}{W}} Z_1 \tag{57}$$

$$N_2 = \sqrt{\frac{-2\log W}{W}} Z_2 \tag{58}$$

$$W = Z_1^2 + Z_2^2 \tag{59}$$

Box-Muller Method in Polar Representation

- Clearly $W = Z_1^2 + Z_2^2$ and $\cos\Theta = \frac{Z_1}{\sqrt{W}}$, and $\sin\Theta = \frac{Z_2}{\sqrt{W}}$;
- Then the Box-Muller transform can be written as

$$N_1 = \sqrt{\frac{-2\log W}{W}} Z_1 \tag{57}$$

$$N_2 = \sqrt{\frac{-2\log W}{W}} Z_2 \tag{58}$$

$$W = Z_1^2 + Z_2^2 \tag{59}$$

Polar Method

To sample from the unit disk, we adopt the truncated rejection method:

- sample (V_1, V_2) uniformly from $[-1, 1] \times [-1, 1]$;
- until $0 < V_1^2 + V_2^2 \leq 1$, set $(Z_1, Z_2) = (V_1, V_2)$.

Polar Method

To sample from the unit disk, we adopt the truncated rejection method:

- sample (V_1, V_2) uniformly from $[-1, 1] \times [-1, 1]$;
- until $0 < V_1^2 + V_2^2 \leq 1$, set $(Z_1, Z_2) = (V_1, V_2)$.

Poissonian noise

- Each image data $I_{i,j}$ is the random sampling from a random variable, which is Poissonian distributed with mean $d_{i,j}$.
- Usually those random variables are independent.
- Then we have

$$\mathbf{Pr}(I|d) = \prod_{i,j} \mathbf{Pr}(I_{i,j}|d_{i,j}) \quad = \prod_{i,j} \frac{{d_{i,j}}^{I_{i,j}}}{I_{i,j}!} \exp\left(-d_{i,j}\right) \tag{60}$$

Poissonian noise

- Each image data $I_{i,j}$ is the random sampling from a random variable, which is Poissonian distributed with mean $d_{i,j}$.
- Usually those random variables are independent.
- Then we have

$$\mathbf{Pr}(I|d) = \prod_{i,j} \mathbf{Pr}(I_{i,j}|d_{i,j}) \quad = \prod_{i,j} \frac{{d_{i,j}}^{I_{i,j}}}{I_{i,j}!} \exp\left(-d_{i,j}\right) \tag{60}$$

Poissonian noise

- Each image data $I_{i,j}$ is the random sampling from a random variable, which is Poissonian distributed with mean $d_{i,j}$.
- Usually those random variables are independent.
- Then we have

$$\mathbf{Pr}(I|d) = \prod_{i,j} \mathbf{Pr}(I_{i,j}|d_{i,j}) \quad = \prod_{i,j} \frac{{d_{i,j}}^{I_{i,j}}}{I_{i,j}!} \exp\left(-d_{i,j}\right) \tag{60}$$

Mixed noise: Poissonian + Gaussian

- In this model, the image data $I_{i,j}$ consist of two parts,

$$I_{i,j} = P_{i,j} + G_{i,j} \quad (61)$$

where

  - $O_{i,j}$ are the data from the object, which are independently Poissonian distributed random variables with mean $d_{i,j}$;
  - $g_{i,j}$ are real-valued random variables accounting for the noise that is present in the device and environment.

Mixed noise: Poissonian + Gaussian

- In this model, the image data $I_{i,j}$ consist of two parts,

$$I_{i,j} = P_{i,j} + G_{i,j} \tag{61}$$

  where
  - $O_{i,j}$ are the data from the object, which are independently Poissonian distributed random variables with mean $d_{i,j}$;
  - $g_{i,j}$ are real-valued random variables accounting for the noise that is present in the device and environment.

Mixed noise: Poissonian + Gaussian

- In this model, the image data $I_{i,j}$ consist of two parts,

$$I_{i,j} = P_{i,j} + G_{i,j} \tag{61}$$

  where
  - $O_{i,j}$ are the data from the object, which are independently Poissonian distributed random variables with mean $d_{i,j}$;
  - $g_{i,j}$ are real-valued random variables accounting for the noise that is present in the device and environment.

Homework: Image Quality and Image Noise

- Use images corrupted with various noises to demonstrated the performance of image quality measures such as MSE (21), SNR (22) and PSNR (23), etc.
- Which measure do you think statitically reasonable?
- You cal also consider to use those measures from [D. R. Wilson and T. R. Martinez, 1997].
- Please refer to [Wang et al., 2004] and references therein.

Homework: Image Quality and Image Noise

- Use images corrupted with various noises to demonstrated the performance of image quality measures such as MSE (21), SNR (22) and PSNR (23), etc.
- Which measure do you think statitically reasonable?
- You cal also consider to use those measures from [D. R. Wilson and T. R. Martinez, 1997].
- Please refer to [Wang et al., 2004] and references therein.

Homework: Image Quality and Image Noise

- Use images corrupted with various noises to demonstrated the performance of image quality measures such as MSE (21), SNR (22) and PSNR (23), etc.
- Which measure do you think statitically reasonable?
- You cal also consider to use those measures from [D. R. Wilson and T. R. Martinez, 1997].
- Please refer to [Wang et al., 2004] and references therein.

Homework: Image Quality and Image Noise

- Use images corrupted with various noises to demonstrated the performance of image quality measures such as MSE (21), SNR (22) and PSNR (23), etc.
- Which measure do you think statitically reasonable?
- You cal also consider to use those measures from [D. R. Wilson and T. R. Martinez, 1997].
- Please refer to [Wang et al., 2004] and references therein.

Color

- Human color perception adds a subjective layer on top of underlying objective physical properties — the wavelength of electromagnetic radiation.
- Consequently, color may be considered a psychophysical phenomenon.
- Color has long been used in painting, photography and films to display the surrounding world to humans in a similar way in which it is perceived.
- There is considerable literature on the variants in the naming of colors across languages — a very subtle affair.

Color

- Human color perception adds a subjective layer on top of underlying objective physical properties — the wavelength of electromagnetic radiation.
- Consequently, color may be considered a psychophysical phenomenon.
- Color has long been used in painting, photography and films to display the surrounding world to humans in a similar way in which it is perceived.
- There is considerable literature on the variants in the naming of colors across languages — a very subtle affair.

Color

- Human color perception adds a subjective layer on top of underlying objective physical properties — the wavelength of electromagnetic radiation.
- Consequently, color may be considered a psychophysical phenomenon.
- Color has long been used in painting, photography and films to display the surrounding world to humans in a similar way in which it is perceived.
- There is considerable literature on the variants in the naming of colors across languages — a very subtle affair.

Color

- Human color perception adds a subjective layer on top of underlying objective physical properties — the wavelength of electromagnetic radiation.
- Consequently, color may be considered a psychophysical phenomenon.
- Color has long been used in painting, photography and films to display the surrounding world to humans in a similar way in which it is perceived.
- There is considerable literature on the variants in the naming of colors across languages — a very subtle affair.

Color Constancy

- The human visual system is not very precise in perceiving- color in absolute terms;
  - if we wish to express our notion of color precisely,
  - we would describe it relative to some widely used color which is used as a standard.
- There are whole industries which present images to humans, and hence a desire for `color constancy`.
- In computer vision, we have the advantage of using a camera as a measuring device, **which yields measurements in absolute quantities**.

Color Constancy

- The human visual system is not very precise in perceiving- color in absolute terms;
  - if we wish to express our notion of color precisely,
  - we would describe it relative to some widely used color which is used as a standard.
- There are whole industries which present images to humans, and hence a desire for `color constancy`.
- In computer vision, we have the advantage of using a camera as a measuring device, **which yields measurements in absolute quantities**.

Color Constancy

- The human visual system is not very precise in perceiving- color in absolute terms;
  - if we wish to express our notion of color precisely,
  - we would describe it relative to some widely used color which is used as a standard.
- There are whole industries which present images to humans, and hence a desire for `color constancy`.
- In computer vision, we have the advantage of using a camera as a measuring device, **which yields measurements in absolute quantities**.

Color Constancy

- The human visual system is not very precise in perceiving- color in absolute terms;
  - if we wish to express our notion of color precisely,
  - we would describe it relative to some widely used color which is used as a standard.
- There are whole industries which present images to humans, and hence a desire for `color constancy`.
- In computer vision, we have the advantage of using a camera as a measuring device, **which yields measurements in absolute quantities**.

Color Constancy

- The human visual system is not very precise in perceiving- color in absolute terms;
  - if we wish to express our notion of color precisely,
  - we would describe it relative to some widely used color which is used as a standard.
- There are whole industries which present images to humans, and hence a desire for `color constancy`.
- In computer vision, we have the advantage of using a camera as a measuring device, **which yields measurements in absolute quantities**.

Newton

- Newton reported in the 17th century that white light from the sun is a spectral mixture,
- and used the optical prism to perform decomposition.
- This was a radical idea to propose at time.
- Over 100 years later influential scientists and philosophers such as Goethe refused to believe it.

Newton

- Newton reported in the 17th century that white light from the sun is a spectral mixture,
- and used the optical prism to perform decomposition.
- This was a radical idea to propose at time.
- Over 100 years later influential scientists and philosophers such as Goethe refused to believe it.

Newton

- Newton reported in the 17th century that white light from the sun is a spectral mixture,
- and used the optical prism to perform decomposition.
- This was a radical idea to propose at time.
- Over 100 years later influential scientists and philosophers such as Goethe refused to believe it.

Newton

- Newton reported in the 17th century that white light from the sun is a spectral mixture,
- and used the optical prism to perform decomposition.
- This was a radical idea to propose at time.
- Over 100 years later influential scientists and philosophers such as Goethe refused to believe it.

Outline

Basic concepts
- Image functions
- Image formation
- Image as a stochastic process

Image digitization
- Sampling
- Quantization

Digital image properties
- Metric and topological properties of digital images
  - Metric properties of digital images
  - Topological properties of digital images
  - Contiguity paradoxes
  - Other topological and geometrical properties
- Histogram
- Entropy
- Visual perception of the image
  - Contrast
  - Acuity
  - Visual Illusions
- Image quality

Electromagnetic Spectrum

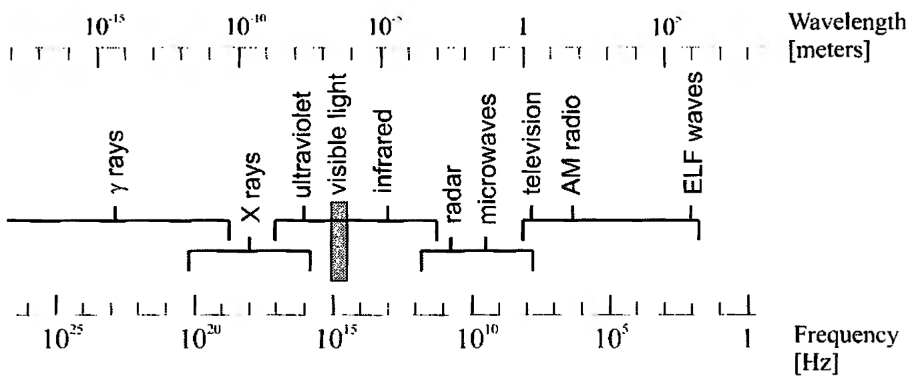

Visual Spectrum

- The electromagnetic spectrum visible to a human, is with the wavelength λ from approximately 380 nm to 740 nm.
- The intensity of irradiation for different wavelengths λ is called the `power spectrum` (or `power spectrum distribution` $S(\lambda)$).
- Visible colors with the wavelengths shown are called spectral colors

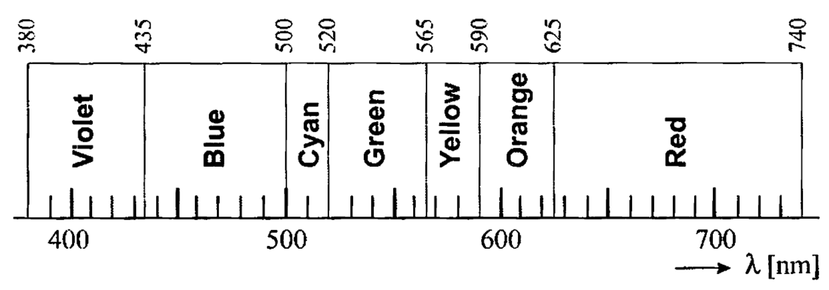

Visual Spectrum

- The electromagnetic spectrum visible to a human, is with the wavelength λ from approximately 380 nm to 740 nm.
- The intensity of irradiation for different wavelengths λ is called the `power spectrum` (or `power spectrum distribution` $S(\lambda)$.
- Visible colors with the wavelengths shown are called spectral colors

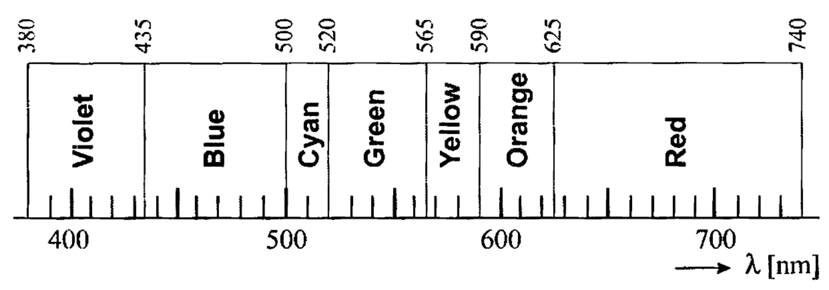

Visual Spectrum

- The electromagnetic spectrum visible to a human, is with the wavelength λ from approximately 380 nm to 740 nm.
- The intensity of irradiation for different wavelengths λ is called the `power spectrum` (or `power spectrum distribution` $S(\lambda)$).
- Visible colors with the wavelengths shown are called spectral colors

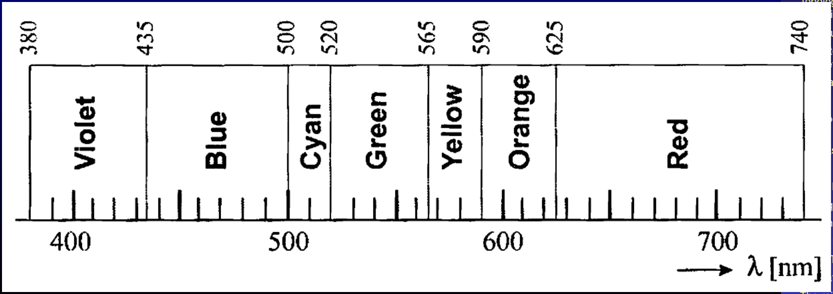

Primary Colors

- Colors can be represented as combinations of the primary colors, e.g., red, green, and blue;
- For the purposes of standardization they have been defined as 700 nm, 546.1 nm, and 435.8 nm, respectively [Pratt, 1978].
- But this standardization does not imply that all colors can be synthesized as combinations of these three.

Primary Colors

- Colors can be represented as combinations of the primary colors, e.g., red, green, and blue;
- For the purposes of standardization they have been defined as 700 nm, 546.1 nm, and 435.8 nm, respectively [Pratt, 1978].
- But this standardization does not imply that all colors can be synthesized as combinations of these three.

Primary Colors

- Colors can be represented as combinations of the primary colors, e.g., red, green, and blue;
- For the purposes of standardization they have been defined as 700 nm, 546.1 nm, and 435.8 nm, respectively [Pratt, 1978].
- But this standardization does not imply that all colors can be synthesized as combinations of these three.

Surface Reflection

- Two predominant physical mechanisms describing what happens when a surface is irradiated.
- First, the surface reflection rebounds incoming energy in a similar way to a mirror.
- The spectrum of the reflected light remains the same as that of the illuminant and it is independent of the surface — recall that shiny metals ’do not have a color’.

Surface Reflection

- Two predominant physical mechanisms describing what happens when a surface is irradiated.
- First, the surface reflection rebounds incoming energy in a similar way to a mirror.
- The spectrum of the reflected light remains the same as that of the illuminant and it is independent of the surface — recall that shiny metals ’do not have a color’.

Surface Reflection

- Two predominant physical mechanisms describing what happens when a surface is irradiated.
- First, the surface reflection rebounds incoming energy in a similar way to a mirror.
- The spectrum of the reflected light remains the same as that of the illuminant and it is independent of the surface — recall that shiny metals ’do not have a color’.

Body Reflection

- Second, the energy diffuses into the material and reflects randomly from the internal pigment in the matter.
- This mechanism is called body reflection and is predominant in dielectrics as plastic or paints.
- Colors are caused by the properties of pigment particles which absorb certain wavelengths from the incoming illuminant wavelength spectrum.

Body Reflection

- Second, the energy diffuses into the material and reflects randomly from the internal pigment in the matter.
- This mechanism is called body reflection and is predominant in dielectrics as plastic or paints.
- Colors are caused by the properties of pigment particles which absorb certain wavelengths from the incoming illuminant wavelength spectrum.

Body Reflection

- Second, the energy diffuses into the material and reflects randomly from the internal pigment in the matter.
- This mechanism is called body reflection and is predominant in dielectrics as plastic or paints.
- Colors are caused by the properties of pigment particles which absorb certain wavelengths from the incoming illuminant wavelength spectrum.

Surface and Body Reflections

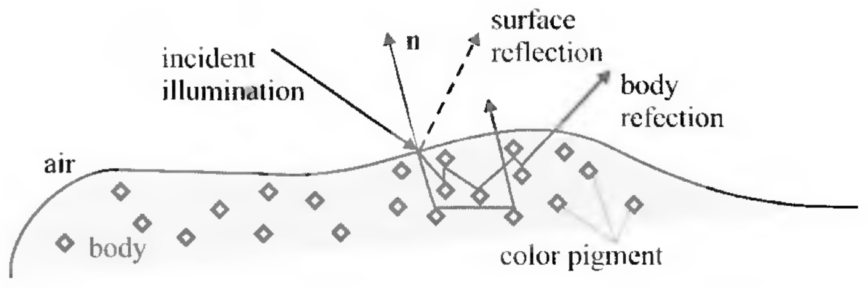

Spectrophotometer

- Most sensors used for color capture, e.g., in cameras, do not have direct access to color;
- the exception is a spectrophotometer which in principle resembles Newton's prism.
- Incoming irradiation is decomposed into spectral colors and intensity along the spectrum is measured in a narrow wavelength band,
  - for instance, by a mechanically moved point sensor.
- Actual spectrophotometers use diffraction gratings instead of a glass prism.

Spectrophotometer

- Most sensors used for color capture, e.g., in cameras, do not have direct access to color;
- the exception is a spectrophotometer which in principle resembles Newton's prism.
- Incoming irradiation is decomposed into spectral colors and intensity along the spectrum is measured in a narrow wavelength band,
  - for instance, by a mechanically moved point sensor.
- Actual spectrophotometers use diffraction gratings instead of a glass prism.

Spectrophotometer

- Most sensors used for color capture, e.g., in cameras, do not have direct access to color;
- the exception is a spectrophotometer which in principle resembles Newton's prism.
- Incoming irradiation is decomposed into spectral colors and intensity along the spectrum is measured in a narrow wavelength band,
  - for instance, by a mechanically moved point sensor.
- Actual spectrophotometers use diffraction gratings instead of a glass prism.

Spectrophotometer

- Most sensors used for color capture, e.g., in cameras, do not have direct access to color;
- the exception is a spectrophotometer which in principle resembles Newton's prism.
- Incoming irradiation is decomposed into spectral colors and intensity along the spectrum is measured in a narrow wavelength band,
  - for instance, by a mechanically moved point sensor.
- Actual spectrophotometers use diffraction gratings instead of a glass prism.

Spectrophotometer

- Most sensors used for color capture, e.g., in cameras, do not have direct access to color;
- the exception is a spectrophotometer which in principle resembles Newton's prism.
- Incoming irradiation is decomposed into spectral colors and intensity along the spectrum is measured in a narrow wavelength band,
  - for instance, by a mechanically moved point sensor.
- Actual spectrophotometers use diffraction gratings instead of a glass prism.

Multispectral images

- Light intensities measured in several narrow bands of wavelengths are collected in a vector describing each pixel.
- Each spectral band is digitized independently and is represented by an individual digital image function as if it were a monochromatic image.
- `Multispectral images` are commonly used in remote sensing from satellites, airborne sensors and in industry.
- Wavelength usually span from ultraviolet through the visible section to infrared.
- Seven or a dozen wavelength bands are common.
- For instance, the LANDSAT 4 satellite transmits digitized images in five spectral bands from near-ultraviolet to infrared.

Multispectral images

- Light intensities measured in several narrow bands of wavelengths are collected in a vector describing each pixel.
- Each spectral band is digitized independently and is represented by an individual digital image function as if it were a monochromatic image.
- `Multispectral images` are commonly used in remote sensing from satellites, airborne sensors and in industry.
- Wavelength usually span from ultraviolet through the visible section to infrared.
- Seven or a dozen wavelength bands are common.
- For instance, the LANDSAT 4 satellite transmits digitized images in five spectral bands from near-ultraviolet to infrared.

Multispectral images

- Light intensities measured in several narrow bands of wavelengths are collected in a vector describing each pixel.
- Each spectral band is digitized independently and is represented by an individual digital image function as if it were a monochromatic image.
- `Multispectral images` are commonly used in remote sensing from satellites, airborne sensors and in industry.
- Wavelength usually span from ultraviolet through the visible section to infrared.
- Seven or a dozen wavelength bands are common.
- For instance, the LANDSAT 4 satellite transmits digitized images in five spectral bands from near-ultraviolet to infrared.

Multispectral images

- Light intensities measured in several narrow bands of wavelengths are collected in a vector describing each pixel.
- Each spectral band is digitized independently and is represented by an individual digital image function as if it were a monochromatic image.
- `Multispectral images` are commonly used in remote sensing from satellites, airborne sensors and in industry.
- Wavelength usually span from ultraviolet through the visible section to infrared.
- Seven or a dozen wavelength bands are common.
- For instance, the LANDSAT 4 satellite transmits digitized images in five spectral bands from near-ultraviolet to infrared.

Multispectral images

- Light intensities measured in several narrow bands of wavelengths are collected in a vector describing each pixel.
- Each spectral band is digitized independently and is represented by an individual digital image function as if it were a monochromatic image.
- `Multispectral images` are commonly used in remote sensing from satellites, airborne sensors and in industry.
- Wavelength usually span from ultraviolet through the visible section to infrared.
- Seven or a dozen wavelength bands are common.
- For instance, the LANDSAT 4 satellite transmits digitized images in five spectral bands from near-ultraviolet to infrared.

Multispectral images

- Light intensities measured in several narrow bands of wavelengths are collected in a vector describing each pixel.
- Each spectral band is digitized independently and is represented by an individual digital image function as if it were a monochromatic image.
- `Multispectral images` are commonly used in remote sensing from satellites, airborne sensors and in industry.
- Wavelength usually span from ultraviolet through the visible section to infrared.
- Seven or a dozen wavelength bands are common.
- For instance, the LANDSAT 4 satellite transmits digitized images in five spectral bands from near-ultraviolet to infrared.

Outline

- Basic concepts
  - Image functions
  - Image formation
  - Image as a stochastic process
- Image digitization
  - Sampling
  - Quantization
- Digital image properties
  - Metric and topological properties of digital images
    - Metric properties of digital images
    - Topological properties of digital images
    - Contiguity paradoxes
    - Other topological and geometrical properties
  - Histogram
  - Entropy
  - Visual perception of the image
    - Contrast
    - Acuity
    - Visual Illusions
  - Image quality

Receptors

- Evolution has developed a mechanism of indirect color sensing in humans and some animals.
- Three types of sensors receptive to the wavelength of incoming irradiation have been established in humans, thus the term trichromacy.
- Color sensitive receptors on the human retina are the cones.
- The other light sensitive receptors on the retina are the rods which are dedicated to sensing monochromatically in low ambient light conditions.

Receptors

- Evolution has developed a mechanism of indirect color sensing in humans and some animals.
- Three types of sensors receptive to the wavelength of incoming irradiation have been established in humans, thus the term trichromacy.
- Color sensitive receptors on the human retina are the cones.
- The other light sensitive receptors on the retina are the rods which are dedicated to sensing monochromatically in low ambient light conditions.

Receptors

- Evolution has developed a mechanism of indirect color sensing in humans and some animals.
- Three types of sensors receptive to the wavelength of incoming irradiation have been established in humans, thus the term trichromacy.
- Color sensitive receptors on the human retina are the cones.
- The other light sensitive receptors on the retina are the rods which are dedicated to sensing monochromatically in low ambient light conditions.

Receptors

- Evolution has developed a mechanism of indirect color sensing in humans and some animals.
- Three types of sensors receptive to the wavelength of incoming irradiation have been established in humans, thus the term trichromacy.
- Color sensitive receptors on the human retina are the cones.
- The other light sensitive receptors on the retina are the rods which are dedicated to sensing monochromatically in low ambient light conditions.

Cones

- Cones are categorized into three types based on the sensed wavelength range
  - S (short) with maximum sensitivity at $\approx 430\,\mathrm{nm}$,
  - M (medium) at $\approx 560\,\mathrm{nm}$,
  - L (long) at $\approx 610\,\mathrm{nm}$.
- Cones S, M, L are occasionally called cones B, G and R, respectively.
- This is slightly misleading. We do not see red solely because an L cone is activated.
- Light with equally distributed wavelength spectrum looks white to a human, and an unbalanced spectrum appears as some shade of color.

Cones

- Cones are categorized into three types based on the sensed wavelength range
  - S (short) with maximum sensitivity at $\approx 430\,\mathrm{nm}$,
  - M (medium) at $\approx 560\,\mathrm{nm}$,
  - L (long) at $\approx 610\,\mathrm{nm}$.
- Cones S, M, L are occasionally called cones B, G and R, respectively.
- This is slightly misleading. We do not see red solely because an L cone is activated.
- Light with equally distributed wavelength spectrum looks white to a human, and an unbalanced spectrum appears as some shade of color.

Cones

- Cones are categorized into three types based on the sensed wavelength range
  - S (short) with maximum sensitivity at $\approx 430\,\mathrm{nm}$,
  - M (medium) at $\approx 560\,\mathrm{nm}$,
  - L (long) at $\approx 610\,\mathrm{nm}$.
- Cones S, M, L are occasionally called cones B, G and R, respectively.
- This is slightly misleading. We do not see red solely because an L cone is activated.
- Light with equally distributed wavelength spectrum looks white to a human, and an unbalanced spectrum appears as some shade of color.

Cones

- Cones are categorized into three types based on the sensed wavelength range
  - S (short) with maximum sensitivity at $\approx 430\,\mathrm{nm}$,
  - M (medium) at $\approx 560\,\mathrm{nm}$,
  - L (long) at $\approx 610\,\mathrm{nm}$.
- Cones S, M, L are occasionally called cones B, G and R, respectively.
- This is slightly misleading. We do not see red solely because an L cone is activated.
- Light with equally distributed wavelength spectrum looks white to a human, and an unbalanced spectrum appears as some shade of color.

Cones

- Cones are categorized into three types based on the sensed wavelength range
  - S (short) with maximum sensitivity at $\approx 430\,\mathrm{nm}$,
  - M (medium) at $\approx 560\,\mathrm{nm}$,
  - L (long) at $\approx 610\,\mathrm{nm}$.
- Cones S, M, L are occasionally called cones B, G and R, respectively.
- This is slightly misleading. We do not see red solely because an L cone is activated.
- Light with equally distributed wavelength spectrum looks white to a human, and an unbalanced spectrum appears as some shade of color.

Cones

- Cones are categorized into three types based on the sensed wavelength range
  - S (short) with maximum sensitivity at $\approx 430\,\mathrm{nm}$,
  - M (medium) at $\approx 560\,\mathrm{nm}$,
  - L (long) at $\approx 610\,\mathrm{nm}$.
- Cones S, M, L are occasionally called cones B, G and R, respectively.
- This is slightly misleading. We do not see red solely because an L cone is activated.
- Light with equally distributed wavelength spectrum looks white to a human, and an unbalanced spectrum appears as some shade of color.

Cones

- Cones are categorized into three types based on the sensed wavelength range
  - S (short) with maximum sensitivity at $\approx 430\,\mathrm{nm}$,
  - M (medium) at $\approx 560\,\mathrm{nm}$,
  - L (long) at $\approx 610\,\mathrm{nm}$.
- Cones S, M, L are occasionally called cones B, G and R, respectively.
- This is slightly misleading. We do not see red solely because an L cone is activated.
- Light with equally distributed wavelength spectrum looks white to a human, and an unbalanced spectrum appears as some shade of color.

Spectral Response

- The reaction of a photoreceptor or output from a sensor in a camera can be modeled mathematically.
- Let i be the specific type of sensor, $i = 1, 2, 3$, (the retinal cone type S, M, L in the human case).
- Let $R_i(\lambda)$ be the spectral sensitivity of the sensor;
- $I(\lambda)$ be the spectral density of the illumination,
- and $S(\lambda)$ describe how the surface patch reflects each wavelength of the illuminating light.
- The spectral response q_i of the i-th sensor, can be modeled by integration over a certain range of wavelengths

$$q_i = \int_{\lambda_1}^{\lambda_2} I(\lambda) R_i(\lambda) S(\lambda) \quad (62)$$

Spectral Response

- The reaction of a photoreceptor or output from a sensor in a camera can be modeled mathematically.
- Let i be the specific type of sensor, $i = 1, 2, 3$, (the retinal cone type S, M, L in the human case).
- Let $R_i(\lambda)$ be the spectral sensitivity of the sensor;
- $I(\lambda)$ be the spectral density of the illumination,
- and $S(\lambda)$ describe how the surface patch reflects each wavelength of the illuminating light.
- The spectral response q_i of the i-th sensor, can be modeled by integration over a certain range of wavelengths

$$q_i = \int_{\lambda_1}^{\lambda_2} I(\lambda) R_i(\lambda) S(\lambda) \qquad (62)$$

Spectral Response

- The reaction of a photoreceptor or output from a sensor in a camera can be modeled mathematically.
- Let i be the specific type of sensor, $i = 1, 2, 3$, (the retinal cone type S, M, L in the human case).
- Let $R_i(\lambda)$ be the spectral sensitivity of the sensor;
- $I(\lambda)$ be the spectral density of the illumination,
- and $S(\lambda)$ describe how the surface patch reflects each wavelength of the illuminating light.
- The spectral response q_i of the i-th sensor, can be modeled by integration over a certain range of wavelengths

$$q_i = \int_{\lambda_1}^{\lambda_2} I(\lambda) R_i(\lambda) S(\lambda) \tag{62}$$

Spectral Response

- The reaction of a photoreceptor or output from a sensor in a camera can be modeled mathematically.
- Let i be the specific type of sensor, $i = 1, 2, 3$, (the retinal cone type S, M, L in the human case).
- Let $R_i(\lambda)$ be the spectral sensitivity of the sensor;
- $I(\lambda)$ be the spectral density of the illumination,
- and $S(\lambda)$ describe how the surface patch reflects each wavelength of the illuminating light.
- The spectral response q_i of the i-th sensor, can be modeled by integration over a certain range of wavelengths

$$q_i = \int_{\lambda_1}^{\lambda_2} I(\lambda) R_i(\lambda) S(\lambda) \tag{62}$$

Spectral Response

- The reaction of a photoreceptor or output from a sensor in a camera can be modeled mathematically.
- Let i be the specific type of sensor, $i = 1, 2, 3$, (the retinal cone type S, M, L in the human case).
- Let $R_i(\lambda)$ be the spectral sensitivity of the sensor;
- $I(\lambda)$ be the spectral density of the illumination,
- and $S(\lambda)$ describe how the surface patch reflects each wavelength of the illuminating light.
- The spectral response q_i of the i-th sensor, can be modeled by integration over a certain range of wavelengths

$$q_i = \int_{\lambda_1}^{\lambda_2} I(\lambda) R_i(\lambda) S(\lambda) \tag{62}$$

Spectral Response

- The reaction of a photoreceptor or output from a sensor in a camera can be modeled mathematically.
- Let i be the specific type of sensor, $i = 1, 2, 3$, (the retinal cone type S, M, L in the human case).
- Let $R_i(\lambda)$ be the spectral sensitivity of the sensor;
- $I(\lambda)$ be the spectral density of the illumination,
- and $S(\lambda)$ describe how the surface patch reflects each wavelength of the illuminating light.
- The spectral response q_i of the i-th sensor, can be modeled by integration over a certain range of wavelengths

$$q_i = \int_{\lambda_1}^{\lambda_2} I(\lambda) R_i(\lambda) S(\lambda) \tag{62}$$

Cones Relative Sensitivity

- Consider the cone types S, M, L.
- Only in the ideal case, when the illumination is perfectly white, i.e., $I(\lambda) = 1$,
- (q_S, q_M, q_L) is an estimate of the color of the surface.
- The figure illustrates qualitatively the relative sensitivities of S, M, L cones.
- Measurements carefully were taken with the white light source at the cornea so that absorption of wavelength in cornea, lens and inner pigments of the eye is taken into account [Wandell, 1995].

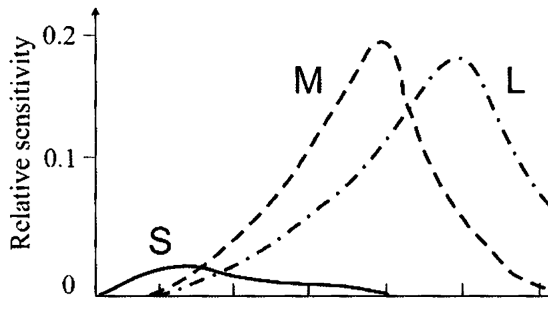

Relative sensitivities of S, M, L cones of the human eye to wavelength.

Cones Relative Sensitivity

- Consider the cone types S, M, L.
- Only in the ideal case, when the illumination is perfectly white, i.e., $I(\lambda) = 1$,
- (q_S, q_M, q_L) is an estimate of the color of the surface.
- The figure illustrates qualitatively the relative sensitivities of S, M, L cones.
- Measurements carefully were taken with the white light source at the cornea so that absorption of wavelength in cornea, lens and inner pigments of the eye is taken into account [Wandell, 1995].

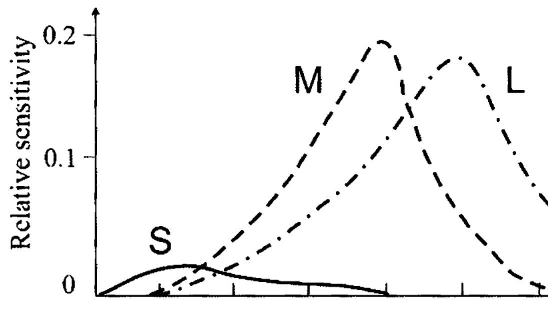

Relative sensitivities of S, M, L cones of the human eye to wavelength.

Cones Relative Sensitivity

- Consider the cone types S, M, L.
- Only in the ideal case, when the illumination is perfectly white, i.e., $I(\lambda) = 1$,
- (q_S, q_M, q_L) is an estimate of the color of the surface.
- The figure illustrates qualitatively the relative sensitivities of S, M, L cones.
- Measurements carefully were taken with the white light source at the cornea so that absorption of wavelength in cornea, lens and inner pigments of the eye is taken into account [Wandell, 1995].

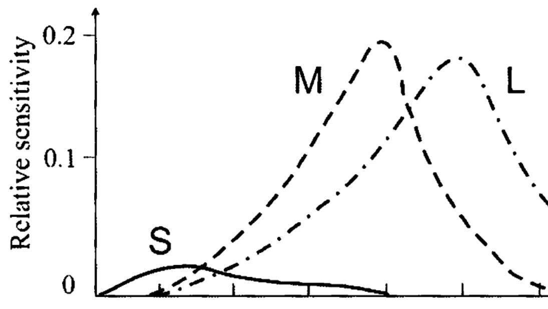

Relative sensitivities of S, M, L cones of the human eye to wavelength.

Cones Relative Sensitivity

- Consider the cone types S, M, L.
- Only in the ideal case, when the illumination is perfectly white, i.e., $I(\lambda) = 1$,
- (q_S, q_M, q_L) is an estimate of the color of the surface.
- The figure illustrates qualitatively the relative sensitivities of S, M, L cones.
- Measurements carefully were taken with the white light source at the cornea so that absorption of wavelength in cornea, lens and inner pigments of the eye is taken into account [Wandell, 1995].

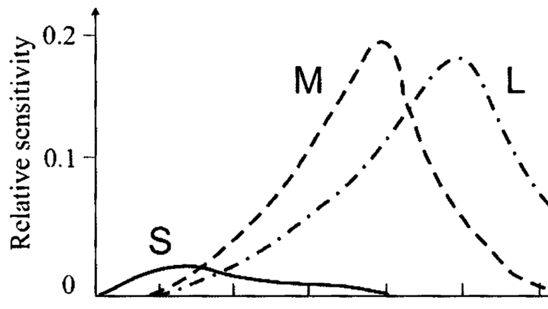

Relative sensitivities of S, M, L cones of the human eye to wavelength.

Cones Relative Sensitivity

- Consider the cone types S, M, L.
- Only in the ideal case, when the illumination is perfectly white, i.e., $I(\lambda) = 1$,
- (q_S, q_M, q_L) is an estimate of the color of the surface.
- The figure illustrates qualitatively the relative sensitivities of S, M, L cones.
- Measurements carefully were taken with the white light source at the cornea so that absorption of wavelength in cornea, lens and inner pigments of the eye is taken into account [Wandell, 1995].

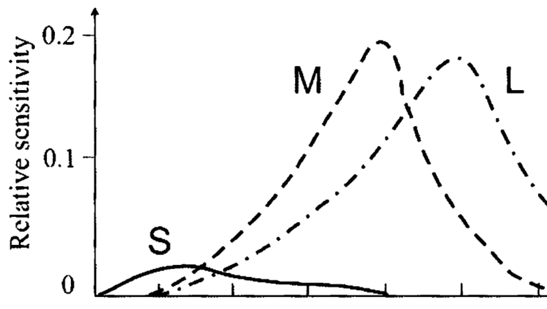

Relative sensitivities of S, M, L cones of the human eye to wavelength.

Cones Relative Sensitivity

- Consider the cone types S, M, L.
- Only in the ideal case, when the illumination is perfectly white, i.e., $I(\lambda) = 1$,
- (q_S, q_M, q_L) is an estimate of the color of the surface.
- The figure illustrates qualitatively the relative sensitivities of S, M, L cones.
- Measurements carefully were taken with the white light source at the cornea so that absorption of wavelength in cornea, lens and inner pigments of the eye is taken into account [Wandell, 1995].

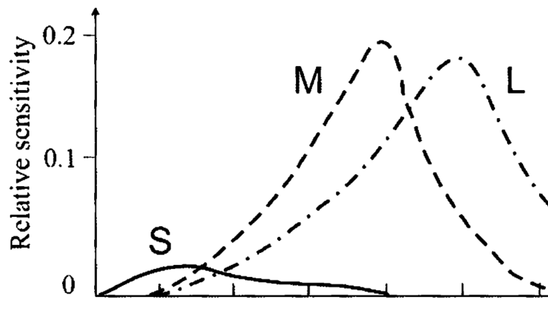

Relative sensitivities of S, M, L cones of the human eye to wavelength.

Cones Relative Sensitivity

- Consider the cone types S, M, L.
- Only in the ideal case, when the illumination is perfectly white, i.e., $I(\lambda) = 1$,
- (q_S, q_M, q_L) is an estimate of the color of the surface.
- The figure illustrates qualitatively the relative sensitivities of S, M, L cones.
- Measurements carefully were taken with the white light source at the cornea so that absorption of wavelength in cornea, lens and inner pigments of the eye is taken into account [Wandell, 1995].

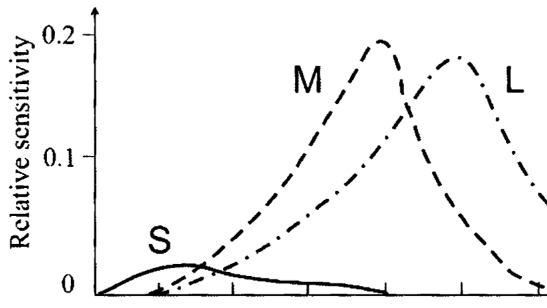

Relative sensitivities of S, M, L cones of the human eye to wavelength.

Cones Relative Sensitivity

- Consider the cone types S, M, L.
- Only in the ideal case, when the illumination is perfectly white, i.e., $I(\lambda) = 1$,
- (q_S, q_M, q_L) is an estimate of the color of the surface.
- The figure illustrates qualitatively the relative sensitivities of S, M, L cones.
- Measurements carefully were taken with the white light source at the cornea so that absorption of wavelength in cornea, lens and inner pigments of the eye is taken into account [Wandell, 1995].

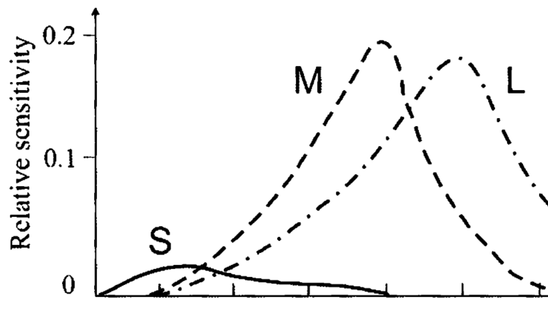

Relative sensitivities of S, M, L cones of the human eye to wavelength.

Color Metamer

- A phenomenon called color metamer is relevant.
- A metamer, in general, means two things that are physically different but perceived as the same.
- Red and green adding to produce yellow is a color mctamer, because yellow could have also been produced by a spectral color.
- The human visual system is fooled into perceiving that red and green is the same as yellow.

Color Metamer

- A phenomenon called color metamer is relevant.
- A metamer, in general, means two things that are physically different but perceived as the same.
- Red and green adding to produce yellow is a color mctamer, because yellow could have also been produced by a spectral color.
- The human visual system is fooled into perceiving that red and green is the same as yellow.

Color Metamer

- A phenomenon called color metamer is relevant.
- A metamer, in general, means two things that are physically different but perceived as the same.
- Red and green adding to produce yellow is a color mctamer, because yellow could have also been produced by a spectral color.
- The human visual system is fooled into perceiving that red and green is the same as yellow.

Color Metamer

- A phenomenon called color metamer is relevant.
- A metamer, in general, means two things that are physically different but perceived as the same.
- Red and green adding to produce yellow is a color mctamer, because yellow could have also been produced by a spectral color.
- The human visual system is fooled into perceiving that red and green is the same as yellow.

Color Matching

- Consider a color matching experiment in which someone is shown a pattern consisting of two adjacent color patches.
- The first patch displays a test light — a spectral color of certain wavelength.
- The second patch is created as an additive combination of three selected primary lights.
- The observer is asked to control the red, green and blue intensities until both patches look identical.
- This color matching experiment is possible because of the color metamer.

Color Matching

- Consider a color matching experiment in which someone is shown a pattern consisting of two adjacent color patches.
- The first patch displays a test light — a spectral color of certain wavelength.
- The second patch is created as an additive combination of three selected primary lights.
- The observer is asked to control the red, green and blue intensities until both patches look identical.
- This color matching experiment is possible because of the color metamer.

Color Matching

- Consider a color matching experiment in which someone is shown a pattern consisting of two adjacent color patches.
- The first patch displays a test light — a spectral color of certain wavelength.
- The second patch is created as an additive combination of three selected primary lights.
- The observer is asked to control the red, green and blue intensities until both patches look identical.
- This color matching experiment is possible because of the color metamer.

Color Matching

- Consider a color matching experiment in which someone is shown a pattern consisting of two adjacent color patches.
- The first patch displays a test light — a spectral color of certain wavelength.
- The second patch is created as an additive combination of three selected primary lights.
- The observer is asked to control the red, green and blue intensities until both patches look identical.
- This color matching experiment is possible because of the color metamer.

Color Matching

- Consider a color matching experiment in which someone is shown a pattern consisting of two adjacent color patches.
- The first patch displays a test light — a spectral color of certain wavelength.
- The second patch is created as an additive combination of three selected primary lights.
- The observer is asked to control the red, green and blue intensities until both patches look identical.
- This color matching experiment is possible because of the color metamer.

Color Matching with Negative Values

- Negative lobes can be seen on the curves for red and green in this figure [Wandell, 1995].
- If the perceptual match has to be obtained then the observer has to add the intensity to the patch corresponding to the spectral color.
- This increase of this intensity is depicted as a decrease in the color matching function. Hence the negative values.

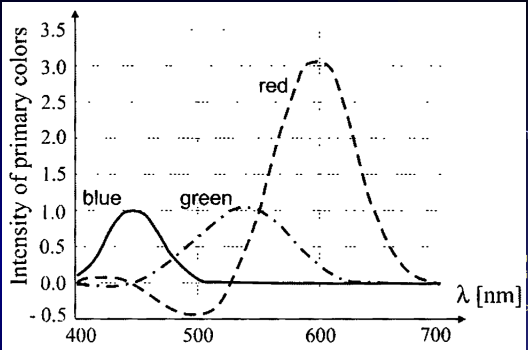

Color matching functions obtained in the color matching experiment. Intensities of the selected primary colors which perceptually match spectral color of given wavelength λ.

Color Matching with Negative Values

- Negative lobes can be seen on the curves for red and green in this figure [Wandell, 1995].
- If the perceptual match has to be obtained then the observer has to add the intensity to the patch corresponding to the spectral color.
- This increase of this intensity is depicted as a decrease in the color matching function. Hence the negative values.

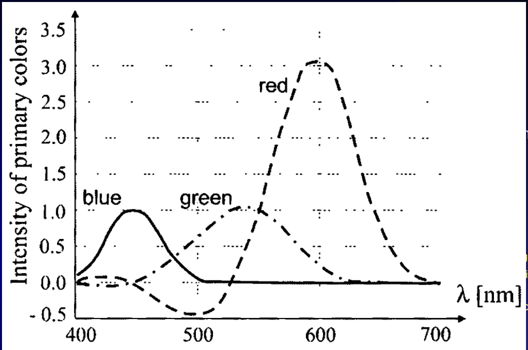

Color matching functions obtained in the color matching experiment. Intensities of the selected primary colors which perceptually match spectral color of given wavelength λ.

Color Matching with Negative Values

- Negative lobes can be seen on the curves for red and green in this figure [Wandell, 1995].
- If the perceptual match has to be obtained then the observer has to add the intensity to the patch corresponding to the spectral color.
- This increase of this intensity is depicted as a decrease in the color matching function. Hence the negative values.

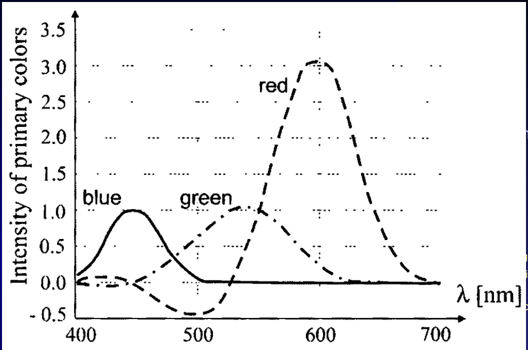

Color matching functions obtained in the color matching experiment. Intensities of the selected primary colors which perceptually match spectral color of given wavelength λ.

Color Matching with Negative Values

- Negative lobes can be seen on the curves for red and green in this figure [Wandell, 1995].
- If the perceptual match has to be obtained then the observer has to add the intensity to the patch corresponding to the spectral color.
- This increase of this intensity is depicted as a decrease in the color matching function. Hence the negative values.

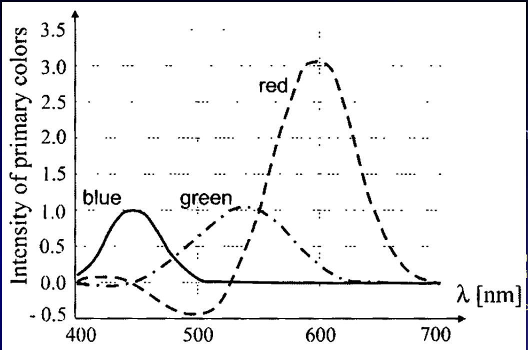

Color matching functions obtained in the color matching experiment. Intensities of the selected primary colors which perceptually match spectral color of given wavelength λ.

Color Matching with Negative Values

- Negative lobes can be seen on the curves for red and green in this figure [Wandell, 1995].
- If the perceptual match has to be obtained then the observer has to add the intensity to the patch corresponding to the spectral color.
- This increase of this intensity is depicted as a decrease in the color matching function. Hence the negative values.

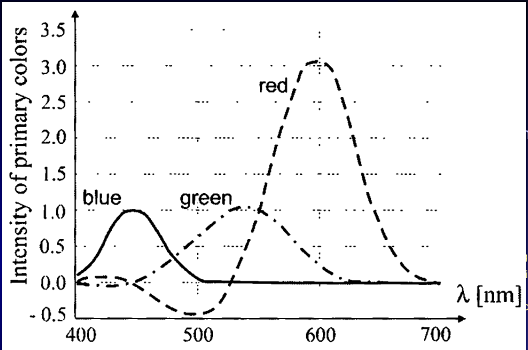

Color matching functions obtained in the color matching experiment. Intensities of the selected primary colors which perceptually match spectral color of given wavelength λ.

Color Matching with Negative Values

- Negative lobes can be seen on the curves for red and green in this figure [Wandell, 1995].
- If the perceptual match has to be obtained then the observer has to add the intensity to the patch corresponding to the spectral color.
- This increase of this intensity is depicted as a decrease in the color matching function. Hence the negative values.

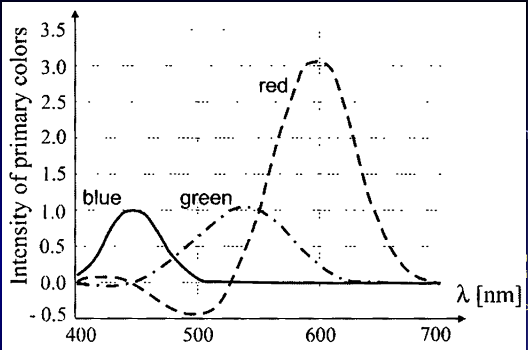

Color matching functions obtained in the color matching experiment. Intensities of the selected primary colors which perceptually match spectral color of given wavelength λ.

Human Vision

- Human vision is prone to various illusions.
- Perceived color is influenced, besides the spectrum of the illuminant, by the colors and scene interpretation surrounding the observed color.
- In addition, eye adaptation to changing light conditions is not very fast and perception is influenced by adaptation.
- Nevertheless, we assume for simplicity that the spectrum of light coming to a point on the retina fully determines the color.

Human Vision

- Human vision is prone to various illusions.
- Perceived color is influenced, besides the spectrum of the illuminant, by the colors and scene interpretation surrounding the observed color.
- In addition, eye adaptation to changing light conditions is not very fast and perception is influenced by adaptation.
- Nevertheless, we assume for simplicity that the spectrum of light coming to a point on the retina fully determines the color.

Human Vision

- Human vision is prone to various illusions.
- Perceived color is influenced, besides the spectrum of the illuminant, by the colors and scene interpretation surrounding the observed color.
- In addition, eye adaptation to changing light conditions is not very fast and perception is influenced by adaptation.
- Nevertheless, we assume for simplicity that the spectrum of light coming to a point on the retina fully determines the color.

Human Vision

- Human vision is prone to various illusions.
- Perceived color is influenced, besides the spectrum of the illuminant, by the colors and scene interpretation surrounding the observed color.
- In addition, eye adaptation to changing light conditions is not very fast and perception is influenced by adaptation.
- Nevertheless, we assume for simplicity that the spectrum of light coming to a point on the retina fully determines the color.

XYZ Color Space

- Color can be defined by almost any set of primaries;
- The world community agreed on primaries and color matching functions which are widely used.
- The color model was introduced as a mathematical abstraction allowing us to express colors as tuples of numbers, typically as three or four values of color components.
- Being motivated by the press and the development of color film, CIE[1] issued a technical standard called XYZ color space in 1931.

[1] International Commission on Illumination, still acting in Lausanne, Switzerland.

XYZ Color Space

- Color can be defined by almost any set of primaries;
- The world community agreed on primaries and color matching functions which are widely used.
- The color model was introduced as a mathematical abstraction allowing us to express colors as tuples of numbers, typically as three or four values of color components.
- Being motivated by the press and the development of color film, CIE[1] issued a technical standard called XYZ color space in 1931.

[1] International Commission on Illumination, still acting in Lausanne, Switzerland.

XYZ Color Space

- Color can be defined by almost any set of primaries;
- The world community agreed on primaries and color matching functions which are widely used.
- The color model was introduced as a mathematical abstraction allowing us to express colors as tuples of numbers, typically as three or four values of color components.
- Being motivated by the press and the development of color film, CIE[1] issued a technical standard called XYZ color space in 1931.

[1] International Commission on Illumination, still acting in Lausanne, Switzerland.

XYZ Color Space

- Color can be defined by almost any set of primaries;
- The world community agreed on primaries and color matching functions which are widely used.
- The color model was introduced as a mathematical abstraction allowing us to express colors as tuples of numbers, typically as three or four values of color components.
- Being motivated by the press and the development of color film, CIE[1] issued a technical standard called XYZ color space in 1931.

[1] International Commission on Illumination, still acting in Lausanne, Switzerland.

CIE Standard

- The standard is given by
  - the three imaginary lights $X = 700.0\,\mathrm{nm}$, $Y = 546.1\,\mathrm{nm}$, $Z = 435.8\,\mathrm{nm}$,
  - the color matching functions $X(\lambda)$, $Y(\lambda)$ and $Z(\lambda)$ corresponding to the perceptual ability of an average human viewing a screen through an aperture providing a $2\,\mathrm{deg}$ field of view.
- The standard is artificial because there is no set of physically realizable primary lights that would yield the color matching functions in the color matching experiment.
- Nevertheless, if we wanted to characterize the imaginary lights then, very roughly speaking, $X \approx$ red, $Y \approx$ green and $Z \approx$ blue.

CIE Standard

- The standard is given by
  - the three imaginary lights $X = 700.0\,\mathrm{nm}$, $Y = 546.1\,\mathrm{nm}$, $Z = 435.8\,\mathrm{nm}$,
  - the color matching functions $X(\lambda)$, $Y(\lambda)$ and $Z(\lambda)$ corresponding to the perceptual ability of an average human viewing a screen through an aperture providing a $2\,\mathrm{deg}$ field of view.
- The standard is artificial because there is no set of physically realizable primary lights that would yield the color matching functions in the color matching experiment.
- Nevertheless, if we wanted to characterize the imaginary lights then, very roughly speaking, $X \approx$ red, $Y \approx$ green and $Z \approx$ blue.

CIE Standard

- The standard is given by
  - the three imaginary lights $X = 700.0\,\mathrm{nm}$, $Y = 546.1\,\mathrm{nm}$, $Z = 435.8\,\mathrm{nm}$,
  - the color matching functions $X(\lambda)$, $Y(\lambda)$ and $Z(\lambda)$ corresponding to the perceptual ability of an average human viewing a screen through an aperture providing a $2\,\mathrm{deg}$ field of view.
- The standard is artificial because there is no set of physically realizable primary lights that would yield the color matching functions in the color matching experiment.
- Nevertheless, if we wanted to characterize the imaginary lights then, very roughly speaking, $X \approx$ red, $Y \approx$ green and $Z \approx$ blue.

CIE Standard

- The standard is given by
  - the three imaginary lights $X = 700.0\,\mathrm{nm}$, $Y = 546.1\,\mathrm{nm}$, $Z = 435.8\,\mathrm{nm}$,
  - the color matching functions $X(\lambda)$, $Y(\lambda)$ and $Z(\lambda)$ corresponding to the perceptual ability of an average human viewing a screen through an aperture providing a $2\,\mathrm{deg}$ field of view.
- The standard is artificial because there is no set of physically realizable primary lights that would yield the color matching functions in the color matching experiment.
- Nevertheless, if we wanted to characterize the imaginary lights then, very roughly speaking, $X \approx$ red, $Y \approx$ green and $Z \approx$ blue.

CIE Standard

- The standard is given by
  - the three imaginary lights $X = 700.0\,\mathrm{nm}$, $Y = 546.1\,\mathrm{nm}$, $Z = 435.8\,\mathrm{nm}$,
  - the color matching functions $X(\lambda)$, $Y(\lambda)$ and $Z(\lambda)$ corresponding to the perceptual ability of an average human viewing a screen through an aperture providing a $2\,\mathrm{deg}$ field of view.
- The standard is artificial because there is no set of physically realizable primary lights that would yield the color matching functions in the color matching experiment.
- Nevertheless, if we wanted to characterize the imaginary lights then, very roughly speaking, $X \approx$ red, $Y \approx$ green and $Z \approx$ blue.

XYZ Standards

- The XYZ color standard fulfills three requirements:
  - A Unlike the color matching experiment yielding negative lobes of color matching functions, the color matching functions of XYZ color space are required to be non-negative;
  - B The value of $Y(\lambda)$ should coincide with the brightness (luminance);
  - C Normalization is performed to assure that the power corresponding to the three color matching functions is equal (i.e., the area under all three curves is equal).

XYZ Standards

- The XYZ color standard fulfills three requirements:
  A Unlike the color matching experiment yielding negative lobes of color matching functions, the color matching functions of XYZ color space are required to be non-negative;
  B The value of $Y(\lambda)$ should coincide with the brightness (luminance);
  C Normalization is performed to assure that the power corresponding to the three color matching functions is equal (i.e., the area under all three curves is equal).

XYZ Standards

- The XYZ color standard fulfills three requirements:
  - A Unlike the color matching experiment yielding negative lobes of color matching functions, the color matching functions of XYZ color space are required to be non-negative;
  - B The value of $Y(\lambda)$ should coincide with the brightness (luminance);
  - C Normalization is performed to assure that the power corresponding to the three color matching functions is equal (i.e., the area under all three curves is equal).

XYZ Standards

- The XYZ color standard fulfills three requirements:
  - A Unlike the color matching experiment yielding negative lobes of color matching functions, the color matching functions of XYZ color space are required to be non-negative;
  - B The value of $Y(\lambda)$ should coincide with the brightness (luminance);
  - C Normalization is performed to assure that the power corresponding to the three color matching functions is equal (i.e., the area under all three curves is equal).

Color Match in XYZ

- A color is a mixture (more precisely a convex combination)

$$c_X X + c_Y Y + c_Z Z \quad (63)$$

where $0 \leq c_X, c_Y, c_Z le 1$, are weights (intensities) in the mixture.

- The resulting color matching functions are shown in the figure [Wandell, 1995].

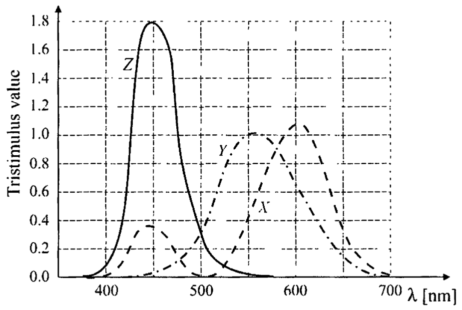

Color Match in XYZ

- A color is a mixture (more precisely a convex combination)

$$c_X X + c_Y Y + c_Z Z \quad (63)$$

where $0 \leq c_X, c_Y, c_Z le 1$, are weights (intensities) in the mixture.

- The resulting color matching functions are shown in the figure [Wandell, 1995].

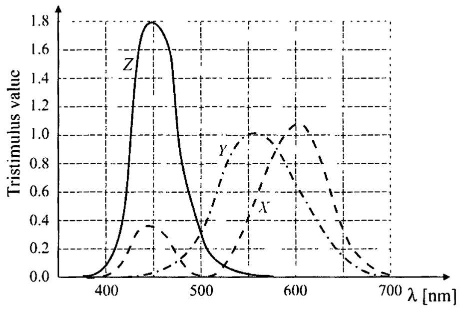

Color Gamut in XYZ

- The subspace of colors perceivable by humans is called the `color gamut` and is demonstrated in the figure.

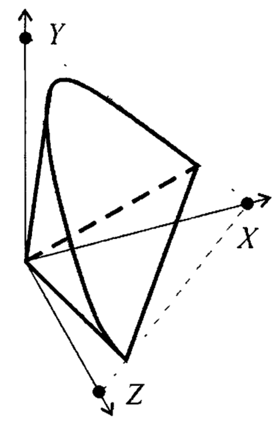

Projection View of Color Gamut in XYZ

- The projection plane is given by the plane passing through extremal points on all three axes, `CIE chromaticity diagram`.
- The new 2D coordinates (x, y) are obtained as

$$x = \frac{X}{X+Y+Z}, \quad (64)$$

$$y = \frac{Y}{X+Y+Z}, \quad (65)$$

$$z = 1 - x - y. \quad (66)$$

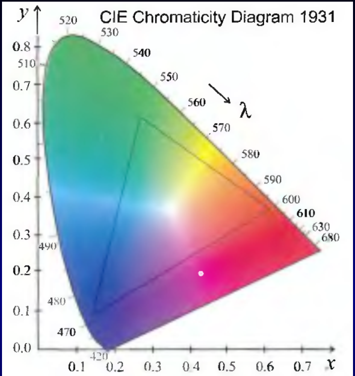

All monochromatic spectra visible to humans map into the curved part of the horseshoe. The triangle depicts a subset of colors spanned by red, green, and blue.

Projection View of Color Gamut in XYZ

- The projection plane is given by the plane passing through extremal points on all three axes, `CIE chromaticity diagram`.
- The new 2D coordinates (x, y) are obtained as

$$x = \frac{X}{X+Y+Z}, \quad (64)$$

$$y = \frac{Y}{X+Y+Z}, \quad (65)$$

$$z = 1 - x - y. \quad (66)$$

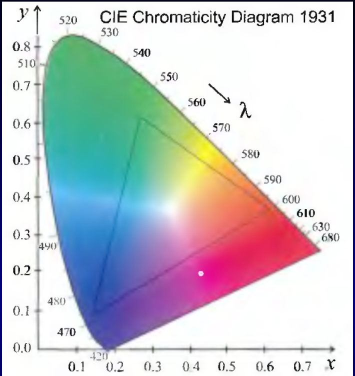

All monochromatic spectra visible to humans map into the curved part of the horseshoe. The triangle depicts a subset of colors spanned by red, green, and blue.

Projection View of Color Gamut in XYZ

- The projection plane is given by the plane passing through extremal points on all three axes, CIE chromaticity diagram.
- The new 2D coordinates (x, y) are obtained as

$$x = \frac{X}{X+Y+Z}, \quad (64)$$

$$y = \frac{Y}{X+Y+Z}, \quad (65)$$

$$z = 1 - x - y. \quad (66)$$

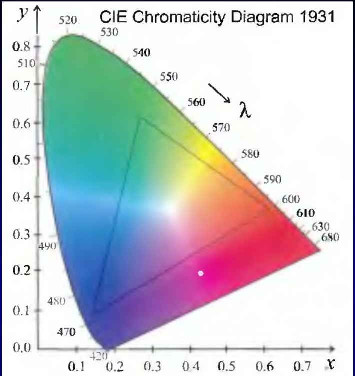

All monochromatic spectra visible to humans map into the curved part of the horseshoe. The triangle depicts a subset of colors spanned by red, green, and blue.

Projection View of Color Gamut in XYZ

- The projection plane is given by the plane passing through extremal points on all three axes, `CIE chromaticity diagram`.
- The new 2D coordinates (x, y) are obtained as

$$x = \frac{X}{X+Y+Z}, \quad (64)$$

$$y = \frac{Y}{X+Y+Z}, \quad (65)$$

$$z = 1 - x - y. \quad (66)$$

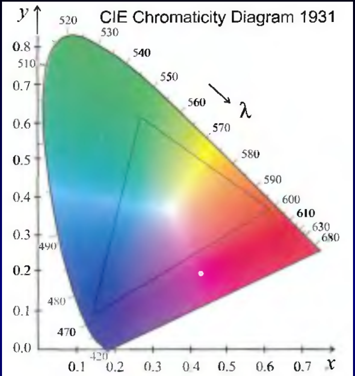

All monochromatic spectra visible to humans map into the curved part of the horseshoe. The triangle depicts a subset of colors spanned by red, green, and blue.

Projection View of Color Gamut in XYZ

- The projection plane is given by the plane passing through extremal points on all three axes, `CIE chromaticity diagram`.
- The new 2D coordinates (x, y) are obtained as

$$x = \frac{X}{X+Y+Z}, \quad (64)$$

$$y = \frac{Y}{X+Y+Z}, \quad (65)$$

$$z = 1 - x - y. \quad (66)$$

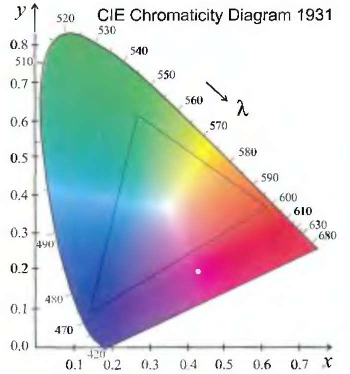

All monochromatic spectra visible to humans map into the curved part of the horseshoe. The triangle depicts a subset of colors spanned by red, green, and blue.

Color Gamuts in XYZ

- Display and printing devices use three selected real primary colors (as opposed to three syntactic primary colors of XYZ color space).
- All possible mixtures of these primary colors fail to cover the whole interior of the horseshoe in CIE chromaticity diagram.
- This situation is demonstrated qualitatively for three particular devices.

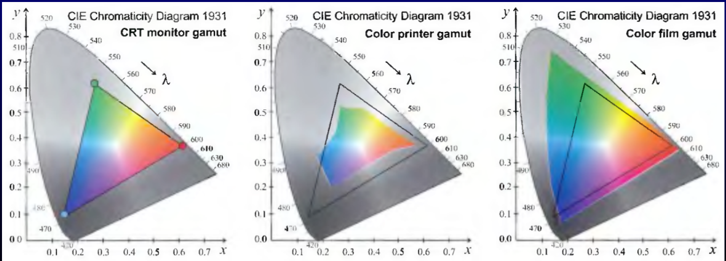

Digital Image Processing
Ming Jiang
Basic concepts
Image functions
Image formation
Image as a stochastic process
Image digitization
Sampling
Quantization
Digital image properties
Metric and topological properties of digital images
Metric properties of digital images
Topological properties of digital images
Contiguity paradoxes
Other topological and geometrical properties
Histogram
Entropy
Visual perception of the image
Contrast
Acuity
Visual Illusions
Image quality
Noise in images
Image Noise
Noise Type
Simulation of noise
The rejection method
The Polar method

Color Gamuts in XYZ

- Display and printing devices use three selected real primary colors (as opposed to three syntactic primary colors of XYZ color space).
- All possible mixtures of these primary colors fail to cover the whole interior of the horseshoe in CIE chromaticity diagram.
- This situation is demonstrated qualitatively for three particular devices.

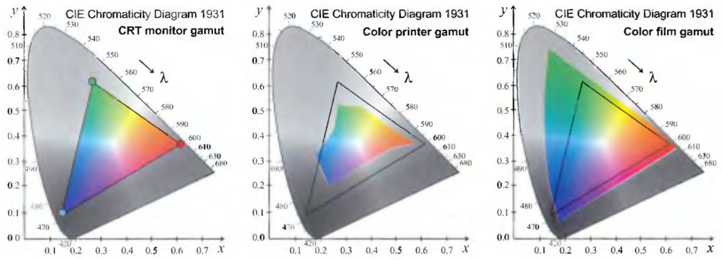

Color Gamuts in XYZ

- Display and printing devices use three selected real primary colors (as opposed to three syntactic primary colors of XYZ color space).
- All possible mixtures of these primary colors fail to cover the whole interior of the horseshoe in CIE chromaticity diagram.
- This situation is demonstrated qualitatively for three particular devices.

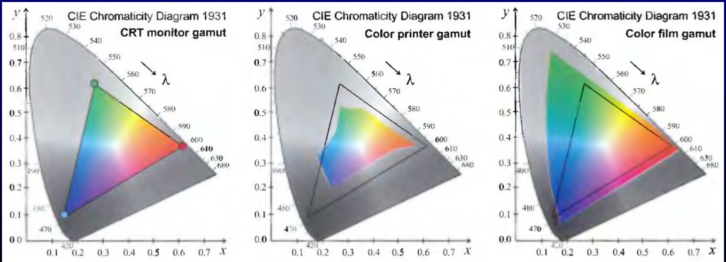

Other Standards

- The CIE standard is an example of an absolute standard, i.e., defining unambiguous representation of color which does not depend on other external factors.
- There are more recent and more precise absolute standards: CIELAB 1976 (ISO 13665) and HunterLab (www.hunterlab.com).
- Later, we will also discuss relative color standards such as RGB color space.
- There are several RGB color spaces used — two computer devices may display the same RGB image differently.

Other Standards

- The CIE standard is an example of an absolute standard, i.e., defining unambiguous representation of color which does not depend on other external factors.
- There are more recent and more precise absolute standards: CIELAB 1976 (ISO 13665) and HunterLab (www.hunterlab.com).
- Later, we will also discuss relative color standards such as RGB color space.
- There are several RGB color spaces used — two computer devices may display the same RGB image differently.

Other Standards

- The CIE standard is an example of an absolute standard, i.e., defining unambiguous representation of color which does not depend on other external factors.
- There are more recent and more precise absolute standards: CIELAB 1976 (ISO 13665) and HunterLab (www.hunterlab.com).
- Later, we will also discuss relative color standards such as RGB color space.
- There are several RGB color spaces used — two computer devices may display the same RGB image differently.

Other Standards

- The CIE standard is an example of an absolute standard, i.e., defining unambiguous representation of color which does not depend on other external factors.
- There are more recent and more precise absolute standards: CIELAB 1976 (ISO 13665) and HunterLab (www.hunterlab.com).
- Later, we will also discuss relative color standards such as RGB color space.
- There are several RGB color spaces used — two computer devices may display the same RGB image differently.

Outline

Color Spaces

- Several different primary colors and corresponding color spaces are used in practice, and these spaces can be transformed into each other.
- If the absolute color space is used then the transformation is the one-to-one mapping and does not lose information (except for rounding errors).
- Because color spaces have their own gamuts, information is lost if the transformed value appears out of the gamut.

Color Spaces

- Several different primary colors and corresponding color spaces are used in practice, and these spaces can be transformed into each other.
- If the absolute color space is used then the transformation is the one-to-one mapping and does not lose information (except for rounding errors).
- Because color spaces have their own gamuts, information is lost if the transformed value appears out of the gamut.

Color Spaces

- Several different primary colors and corresponding color spaces are used in practice, and these spaces can be transformed into each other.
- If the absolute color space is used then the transformation is the one-to-one mapping and does not lose information (except for rounding errors).
- Because color spaces have their own gamuts, information is lost if the transformed value appears out of the gamut.

RGB Color Space

- The RGB color space has its origin in color television where Cathode Ray Tubes (CRT) were used.
- RGB color space is an example of a relative color standard (as opposed to the absolute one, e.g., CIE 1931).
- The primary colors (R-red, G-green and B-blue) mimicked phosphor in CRT luminophore.
- The RGB model uses additive color mixing to inform what kind of light needs to be emitted to produce a given color.
- The value of a particular color is expressed as a vector of three elements — intensities of three primary colors.
- A transformation to a different color space is expressed by a transformation by a 3×3 matrix.

RGB Color Space

- The RGB color space has its origin in color television where Cathode Ray Tubes (CRT) were used.
- RGB color space is an example of a relative color standard (as opposed to the absolute one, e.g., CIE 1931).
- The primary colors (R-red, G-green and B-blue) mimicked phosphor in CRT luminophore.
- The RGB model uses additive color mixing to inform what kind of light needs to be emitted to produce a given color.
- The value of a particular color is expressed as a vector of three elements — intensities of three primary colors.
- A transformation to a different color space is expressed by a transformation by a 3×3 matrix.

RGB Color Space

- The RGB color space has its origin in color television where Cathode Ray Tubes (CRT) were used.
- RGB color space is an example of a relative color standard (as opposed to the absolute one, e.g., CIE 1931).
- The primary colors (R-red, G-green and B-blue) mimicked phosphor in CRT luminophore.
- The RGB model uses additive color mixing to inform what kind of light needs to be emitted to produce a given color.
- The value of a particular color is expressed as a vector of three elements — intensities of three primary colors.
- A transformation to a different color space is expressed by a transformation by a 3×3 matrix.

RGB Color Space

- The RGB color space has its origin in color television where Cathode Ray Tubes (CRT) were used.
- RGB color space is an example of a relative color standard (as opposed to the absolute one, e.g., CIE 1931).
- The primary colors (R-red, G-green and B-blue) mimicked phosphor in CRT luminophore.
- The RGB model uses additive color mixing to inform what kind of light needs to be emitted to produce a given color.
- The value of a particular color is expressed as a vector of three elements — intensities of three primary colors.
- A transformation to a different color space is expressed by a transformation by a 3×3 matrix.

RGB Color Space

- The RGB color space has its origin in color television where Cathode Ray Tubes (CRT) were used.
- RGB color space is an example of a relative color standard (as opposed to the absolute one, e.g., CIE 1931).
- The primary colors (R-red, G-green and B-blue) mimicked phosphor in CRT luminophore.
- The RGB model uses additive color mixing to inform what kind of light needs to be emitted to produce a given color.
- The value of a particular color is expressed as a vector of three elements — intensities of three primary colors.
- A transformation to a different color space is expressed by a transformation by a 3×3 matrix.

RGB Color Space

- The RGB color space has its origin in color television where Cathode Ray Tubes (CRT) were used.
- RGB color space is an example of a relative color standard (as opposed to the absolute one, e.g., CIE 1931).
- The primary colors (R-red, G-green and B-blue) mimicked phosphor in CRT luminophore.
- The RGB model uses additive color mixing to inform what kind of light needs to be emitted to produce a given color.
- The value of a particular color is expressed as a vector of three elements — intensities of three primary colors.
- A transformation to a different color space is expressed by a transformation by a 3×3 matrix.

RGB Representation

- Assume that values for each primary are quantized to $m = 2^n$ values;
- let the highest intensity value be $k = m - 1$;
- then $(0, 0, 0)$ is black, (k, k, k) is (television) white, $(k, 0, 0)$ is 'pure' red, and so on.
- The value $k = 255 = 2^8 - 1$ is common, i.e., 8 bits per color channel.
- There are $256^3 = 2^{24} = 16, 777, 216$ possible colors in such a discretized space.

RGB Representation

- Assume that values for each primary are quantized to $m = 2^n$ values;
- let the highest intensity value be $k = m - 1$;
- then $(0, 0, 0)$ is black, (k, k, k) is (television) white, $(k, 0, 0)$ is 'pure' red, and so on.
- The value $k = 255 = 2^8 - 1$ is common, i.e., 8 bits per color channel.
- There are $256^3 = 2^{24} = 16, 777, 216$ possible colors in such a discretized space.

RGB Representation

- Assume that values for each primary are quantized to $m = 2^n$ values;
- let the highest intensity value be $k = m - 1$;
- then $(0, 0, 0)$ is black, (k, k, k) is (television) white, $(k, 0, 0)$ is 'pure' red, and so on.
- The value $k = 255 = 2^8 - 1$ is common, i.e., 8 bits per color channel.
- There are $256^3 = 2^{24} = 16,777,216$ possible colors in such a discretized space.

RGB Representation

- Assume that values for each primary are quantized to $m = 2^n$ values;
- let the highest intensity value be $k = m - 1$;
- then $(0, 0, 0)$ is black, (k, k, k) is (television) white, $(k, 0, 0)$ is 'pure' red, and so on.
- The value $k = 255 = 2^8 - 1$ is common, i.e., 8 bits per color channel.
- There are $256^3 = 2^{24} = 16,777,216$ possible colors in such a discretized space.

RGB Representation

- Assume that values for each primary are quantized to $m = 2^n$ values;
- let the highest intensity value be $k = m - 1$;
- then $(0, 0, 0)$ is black, (k, k, k) is (television) white, $(k, 0, 0)$ is 'pure' red, and so on.
- The value $k = 255 = 2^8 - 1$ is common, i.e., 8 bits per color channel.
- There are $256^3 = 2^{24} = 16,777,216$ possible colors in such a discretized space.

RGB Color Space Diagram

- RGB color space with primary colors red, green, blue and secondary colors yellow, cyan, magenta.
- Gray-scale images with all intensities lie along the dashed line connecting black and white colors in RGB color space.
- The RGB model may be thought of as a 3D co-ordinatization of color space; note the secondary colors which are combinations of two pure primaries.

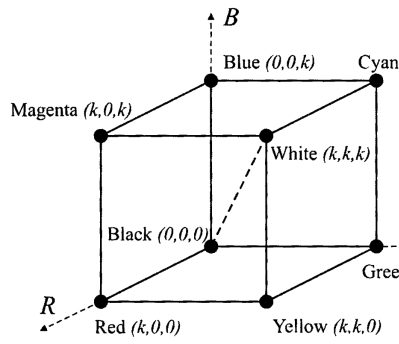

RGB Color Space Diagram

- RGB color space with primary colors red, green, blue and secondary colors yellow, cyan, magenta.
- Gray-scale images with all intensities lie along the dashed line connecting black and white colors in RGB color space.
- The RGB model may be thought of as a 3D co-ordinatization of color space; note the secondary colors which are combinations of two pure primaries.

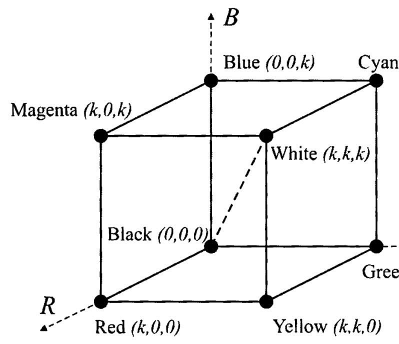

RGB Color Space Diagram

- RGB color space with primary colors red, green, blue and secondary colors yellow, cyan, magenta.
- Gray-scale images with all intensities lie along the dashed line connecting black and white colors in RGB color space.
- The RGB model may be thought of as a 3D co-ordinatization of color space; note the secondary colors which are combinations of two pure primaries.

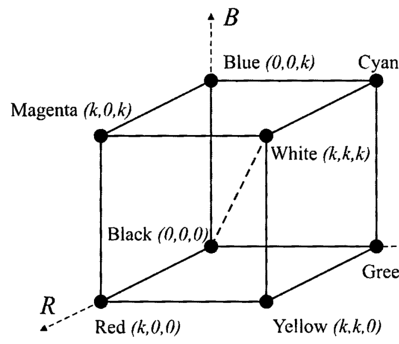

RGB Color Space Diagram

- RGB color space with primary colors red, green, blue and secondary colors yellow, cyan, magenta.
- Gray-scale images with all intensities lie along the dashed line connecting black and white colors in RGB color space.
- The RGB model may be thought of as a 3D co-ordinatization of color space; note the secondary colors which are combinations of two pure primaries.

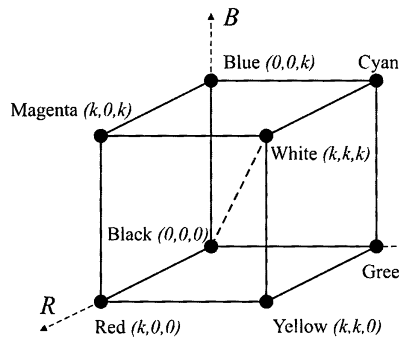

RGB Color Space Diagram

- RGB color space with primary colors red, green, blue and secondary colors yellow, cyan, magenta.
- Gray-scale images with all intensities lie along the dashed line connecting black and white colors in RGB color space.
- The RGB model may be thought of as a 3D co-ordinatization of color space; note the secondary colors which are combinations of two pure primaries.

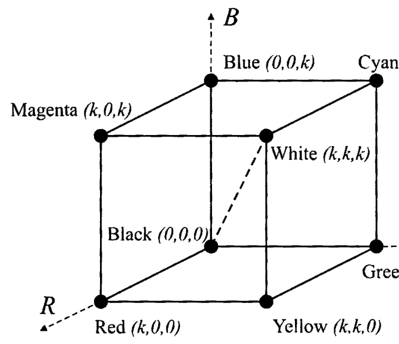

RGB Color Space Diagram

- RGB color space with primary colors red, green, blue and secondary colors yellow, cyan, magenta.
- Gray-scale images with all intensities lie along the dashed line connecting black and white colors in RGB color space.
- The RGB model may be thought of as a 3D co-ordinatization of color space; note the secondary colors which are combinations of two pure primaries.

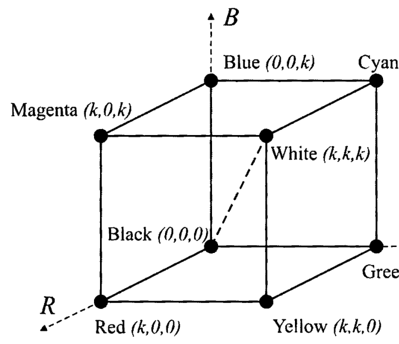

RGB vs XYZ

- There are specific instances of the RGB color model as sRGB, Adobe RGB and Adobe Wide Gamut RGB.
- They differ slightly in transformation matrices and the gamut.
- One of transformations between RGB and XYZ color spaces is

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.24 & -1.54 & -0.50 \\ -0.98 & 1.88 & 0.04 \\ 0.06 & -0.20 & 1.06 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix},$$

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.41 & 0.36 & 0.18 \\ 0.21 & 0.72 & 0.07 \\ 0.02 & 0.12 & 0.95 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}.$$

RGB vs XYZ

- There are specific instances of the RGB color model as sRGB, Adobe RGB and Adobe Wide Gamut RGB.
- They differ slightly in transformation matrices and the gamut.
- One of transformations between RGB and XYZ color spaces is

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.24 & -1.54 & -0.50 \\ -0.98 & 1.88 & 0.04 \\ 0.06 & -0.20 & 1.06 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix},$$

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.41 & 0.36 & 0.18 \\ 0.21 & 0.72 & 0.07 \\ 0.02 & 0.12 & 0.95 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}.$$

RGB vs XYZ

- There are specific instances of the RGB color model as sRGB, Adobe RGB and Adobe Wide Gamut RGB.
- They differ slightly in transformation matrices and the gamut.
- One of transformations between RGB and XYZ color spaces is

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.24 & -1.54 & -0.50 \\ -0.98 & 1.88 & 0.04 \\ 0.06 & -0.20 & 1.06 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix},$$

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.41 & 0.36 & 0.18 \\ 0.21 & 0.72 & 0.07 \\ 0.02 & 0.12 & 0.95 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}.$$

RGB vs XYZ

- There are specific instances of the RGB color model as sRGB, Adobe RGB and Adobe Wide Gamut RGB.
- They differ slightly in transformation matrices and the gamut.
- One of transformations between RGB and XYZ color spaces is

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.24 & -1.54 & -0.50 \\ -0.98 & 1.88 & 0.04 \\ 0.06 & -0.20 & 1.06 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix},$$

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.41 & 0.36 & 0.18 \\ 0.21 & 0.72 & 0.07 \\ 0.02 & 0.12 & 0.95 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}.$$

RGB vs XYZ

- There are specific instances of the RGB color model as sRGB, Adobe RGB and Adobe Wide Gamut RGB.
- They differ slightly in transformation matrices and the gamut.
- One of transformations between RGB and XYZ color spaces is

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.24 & -1.54 & -0.50 \\ -0.98 & 1.88 & 0.04 \\ 0.06 & -0.20 & 1.06 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix},$$

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.41 & 0.36 & 0.18 \\ 0.21 & 0.72 & 0.07 \\ 0.02 & 0.12 & 0.95 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}.$$

RGB vs XYZ

- There are specific instances of the RGB color model as sRGB, Adobe RGB and Adobe Wide Gamut RGB.
- They differ slightly in transformation matrices and the gamut.
- One of transformations between RGB and XYZ color spaces is

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.24 & -1.54 & -0.50 \\ -0.98 & 1.88 & 0.04 \\ 0.06 & -0.20 & 1.06 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix},$$

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.41 & 0.36 & 0.18 \\ 0.21 & 0.72 & 0.07 \\ 0.02 & 0.12 & 0.95 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}.$$

YIQ Color Space

- The US and Japanese color television formerly used YIQ color space.
- The Y component describes intensity
- and I, Q represent colors, corresponding approximately to the amounts of blue and red in the color.
- YIQ is another example of additive color mixing.

YIQ Color Space

- The US and Japanese color television formerly used YIQ color space.
- The Y component describes intensity
- and I, Q represent colors, corresponding approximately to the amounts of blue and red in the color.
- YIQ is another example of additive color mixing.

YIQ Color Space

- The US and Japanese color television formerly used YIQ color space.
- The Y component describes intensity
- and I, Q represent colors, corresponding approximately to the amounts of blue and red in the color.
- YIQ is another example of additive color mixing.

YIQ Color Space

- The US and Japanese color television formerly used YIQ color space.
- The Y component describes intensity
- and I, Q represent colors, corresponding approximately to the amounts of blue and red in the color.
- YIQ is another example of additive color mixing.

YUV Color Space

- This color space corresponds closely to the YUV color model in the PAL television norm.
- YIQ color space is rotated 33 $\deg$ with respect to the YUV color space.
- The YIQ color model is useful since the Y component provides all that is necessary for a monochrome display;
- It exploits advantageous properties of the human visual system, in particular our sensitivity to luminance.

YUV Color Space

- This color space corresponds closely to the YUV color model in the PAL television norm.
- YIQ color space is rotated 33 $\deg$ with respect to the YUV color space.
- The YIQ color model is useful since the Y component provides all that is necessary for a monochrome display;
- It exploits advantageous properties of the human visual system, in particular our sensitivity to luminance.

YUV Color Space

- This color space corresponds closely to the YUV color model in the PAL television norm.
- YIQ color space is rotated 33 $\deg$ with respect to the YUV color space.
- The YIQ color model is useful since the Y component provides all that is necessary for a monochrome display;
- It exploits advantageous properties of the human visual system, in particular our sensitivity to luminance.

YUV Color Space

- This color space corresponds closely to the YUV color model in the PAL television norm.
- YIQ color space is rotated 33 $\deg$ with respect to the YUV color space.
- The YIQ color model is useful since the Y component provides all that is necessary for a monochrome display;
- It exploits advantageous properties of the human visual system, in particular our sensitivity to luminance.

CMY Color Space

- The CMY — for Cyan, Magenta, Yellow — color model uses subtractive color mixing which is used in printing processes.
- It describes what kind of inks need to be applied so the light reflected from the white substrate (paper, painter's canvas) and passing through the inks produces a given color.
- CMYK stores ink values for black in addition.
- Black color can be generated from C, M, Y components.
- As it is abundant in printed documents, it is of advantage to have a special black ink.
- Many CMYK colors spaces are used for different sets of inks, substrates, and press characteristics (which change the color transfer function for each ink and thus change the appearance).

CMY Color Space

- The CMY — for Cyan, Magenta, Yellow — color model uses subtractive color mixing which is used in printing processes.
- It describes what kind of inks need to be applied so the light reflected from the white substrate (paper, painter's canvas) and passing through the inks produces a given color.
- CMYK stores ink values for black in addition.
- Black color can be generated from C, M, Y components.
- As it is abundant in printed documents, it is of advantage to have a special black ink.
- Many CMYK colors spaces are used for different sets of inks, substrates, and press characteristics (which change the color transfer function for each ink and thus change the appearance).

CMY Color Space

- The CMY — for Cyan, Magenta, Yellow — color model uses subtractive color mixing which is used in printing processes.
- It describes what kind of inks need to be applied so the light reflected from the white substrate (paper, painter's canvas) and passing through the inks produces a given color.
- CMYK stores ink values for black in addition.
- Black color can be generated from C, M, Y components.
- As it is abundant in printed documents, it is of advantage to have a special black ink.
- Many CMYK colors spaces are used for different sets of inks, substrates, and press characteristics (which change the color transfer function for each ink and thus change the appearance).

CMY Color Space

- The CMY — for Cyan, Magenta, Yellow — color model uses subtractive color mixing which is used in printing processes.
- It describes what kind of inks need to be applied so the light reflected from the white substrate (paper, painter's canvas) and passing through the inks produces a given color.
- CMYK stores ink values for black in addition.
- Black color can be generated from C, M, Y components.
- As it is abundant in printed documents, it is of advantage to have a special black ink.
- Many CMYK colors spaces are used for different sets of inks, substrates, and press characteristics (which change the color transfer function for each ink and thus change the appearance).

CMY Color Space

- The CMY — for Cyan, Magenta, Yellow — color model uses subtractive color mixing which is used in printing processes.
- It describes what kind of inks need to be applied so the light reflected from the white substrate (paper, painter's canvas) and passing through the inks produces a given color.
- CMYK stores ink values for black in addition.
- Black color can be generated from C, M, Y components.
- As it is abundant in printed documents, it is of advantage to have a special black ink.
- Many CMYK colors spaces are used for different sets of inks, substrates, and press characteristics (which change the color transfer function for each ink and thus change the appearance).

CMY Color Space

- The CMY — for Cyan, Magenta, Yellow — color model uses subtractive color mixing which is used in printing processes.
- It describes what kind of inks need to be applied so the light reflected from the white substrate (paper, painter's canvas) and passing through the inks produces a given color.
- CMYK stores ink values for black in addition.
- Black color can be generated from C, M, Y components.
- As it is abundant in printed documents, it is of advantage to have a special black ink.
- Many CMYK colors spaces are used for different sets of inks, substrates, and press characteristics (which change the color transfer function for each ink and thus change the appearance).

HSV Color Space

- HSV — Hue, Saturation, and Value (also known as HSB, hue, saturation, brightness) is often used by painters because it is closer to their thinking and technique.
- Artists commonly use three to four dozen colors (characterized by the hue; technically, the dominant wavelength).
- If another color is to be obtained then it is mixed from the given ones.
- The painter also wants colors of different saturation.
- E.g., to change 'fire brigade red' to pink, she will mix the 'fire brigade red' with white (andor black) to obtain the desired lower saturation.

HSV Color Space

- HSV — Hue, Saturation, and Value (also known as HSB, hue, saturation, brightness) is often used by painters because it is closer to their thinking and technique.
- Artists commonly use three to four dozen colors (characterized by the hue; technically, the dominant wavelength).
- If another color is to be obtained then it is mixed from the given ones.
- The painter also wants colors of different saturation.
- E.g., to change ’fire brigade red’ to pink, she will mix the ’fire brigade red’ with white (andor black) to obtain the desired lower saturation.

HSV Color Space

- HSV — Hue, Saturation, and Value (also known as HSB, hue, saturation, brightness) is often used by painters because it is closer to their thinking and technique.
- Artists commonly use three to four dozen colors (characterized by the hue; technically, the dominant wavelength).
- If another color is to be obtained then it is mixed from the given ones.
- The painter also wants colors of different saturation.
- E.g., to change ’fire brigade red’ to pink, she will mix the ’fire brigade red’ with white (andor black) to obtain the desired lower saturation.

HSV Color Space

- HSV — Hue, Saturation, and Value (also known as HSB, hue, saturation, brightness) is often used by painters because it is closer to their thinking and technique.
- Artists commonly use three to four dozen colors (characterized by the hue; technically, the dominant wavelength).
- If another color is to be obtained then it is mixed from the given ones.
- The painter also wants colors of different saturation.
- E.g., to change 'fire brigade red' to pink, she will mix the 'fire brigade red' with white (andor black) to obtain the desired lower saturation.

HSV Color Space

- HSV — Hue, Saturation, and Value (also known as HSB, hue, saturation, brightness) is often used by painters because it is closer to their thinking and technique.
- Artists commonly use three to four dozen colors (characterized by the hue; technically, the dominant wavelength).
- If another color is to be obtained then it is mixed from the given ones.
- The painter also wants colors of different saturation.
- E.g., to change ’fire brigade red’ to pink, she will mix the ’fire brigade red’ with white (andor black) to obtain the desired lower saturation.

HSV Color Model

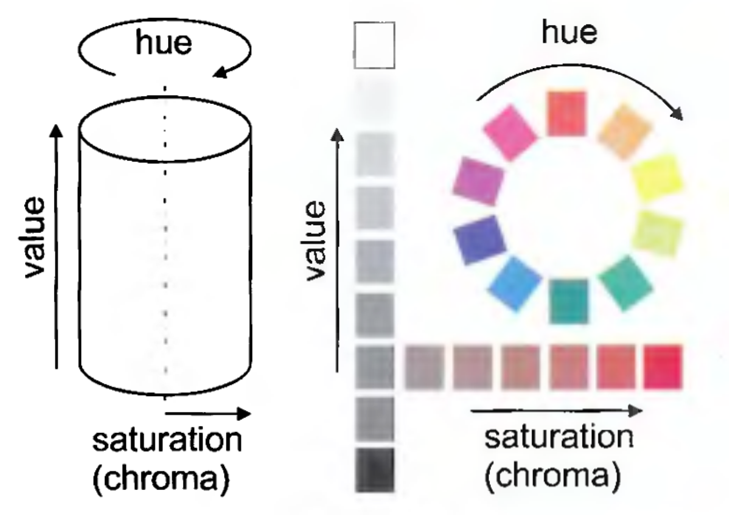

HSV, HSL, HSI

- HSV decouples intensity information from color, while hue and saturation correspond to human perception.
- This representation is very useful for developing image processing algorithms.
- This will become clearer as we proceed to describe image enhancement algorithms.
- HSL (hue, saturation, lightness/luminance), also known as HLS or HSI (hue, saturation, intensity) is similar to HSV.
- ’Lightness’ replaces ’brightness’.
- **The difference is that the brightness of a pure color is equal to the brightness of white, while the lightness of a pure color is equal to the lightness of a medium gray.**

HSV, HSL, HSI

- HSV decouples intensity information from color, while hue and saturation correspond to human perception.
- This representation is very useful for developing image processing algorithms.
- This will become clearer as we proceed to describe image enhancement algorithms.
- HSL (hue, saturation, lightness/luminance), also known as HLS or HSI (hue, saturation, intensity) is similar to HSV.
- 'Lightness' replaces 'brightness'.
- **The difference is that the brightness of a pure color is equal to the brightness of white, while the lightness of a pure color is equal to the lightness of a medium gray.**

HSV, HSL, HSI

- HSV decouples intensity information from color, while hue and saturation correspond to human perception.
- This representation is very useful for developing image processing algorithms.
- This will become clearer as we proceed to describe image enhancement algorithms.
- HSL (hue, saturation, lightness/luminance), also known as HLS or HSI (hue, saturation, intensity) is similar to HSV.
- 'Lightness' replaces 'brightness'.
- **The difference is that the brightness of a pure color is equal to the brightness of white, while the lightness of a pure color is equal to the lightness of a medium gray.**

HSV, HSL, HSI

- HSV decouples intensity information from color, while hue and saturation correspond to human perception.
- This representation is very useful for developing image processing algorithms.
- This will become clearer as we proceed to describe image enhancement algorithms.
- HSL (hue, saturation, lightness/luminance), also known as HLS or HSI (hue, saturation, intensity) is similar to HSV.
- 'Lightness' replaces 'brightness'.
- **The difference is that the brightness of a pure color is equal to the brightness of white, while the lightness of a pure color is equal to the lightness of a medium gray.**

HSV, HSL, HSI

- HSV decouples intensity information from color, while hue and saturation correspond to human perception.
- This representation is very useful for developing image processing algorithms.
- This will become clearer as we proceed to describe image enhancement algorithms.
- HSL (hue, saturation, lightness/luminance), also known as HLS or HSI (hue, saturation, intensity) is similar to HSV.
- ’Lightness’ replaces ’brightness’.
- The difference is that the brightness of a pure color is equal to the brightness of white, while the lightness of a pure color is equal to the lightness of a medium gray.

HSV, HSL, HSI

- HSV decouples intensity information from color, while hue and saturation correspond to human perception.
- This representation is very useful for developing image processing algorithms.
- This will become clearer as we proceed to describe image enhancement algorithms.
- HSL (hue, saturation, lightness/luminance), also known as HLS or HSI (hue, saturation, intensity) is similar to HSV.
- ’Lightness’ replaces ’brightness’.
- **The difference is that the brightness of a pure color is equal to the brightness of white, while the lightness of a pure color is equal to the lightness of a medium gray.**

Summary

Models	Color spaces	Applications
Colorimetric	XYZ	Colorimetric calculations
Device oriented, nonuniform spaces	RGB, UIQ	Storage, processing, coding, color TV
Device oriented, Uniform spaces	LAB, LUV	Color difference, analysis
User oriented	HSL, HSI	Color perception, computer graphics

- Be careful for formulas in books.

Multi-spectral image demonstration

- Multi-spectral image captured by Land-sat TM-5.
- The example uses the images from its 3rd, 4th and 5th spectral bands.
- The matlab script for this example is multi_spectral_image.m.

Multi-spectral image demonstration

- Multi-spectral image captured by Land-sat TM-5.
- The example uses the images from its 3rd, 4th and 5th spectral bands.
- The matlab script for this example is multi_spectral_image.m.

Multi-spectral image demonstration

- Multi-spectral image captured by Land-sat TM-5.
- The example uses the images from its 3rd, 4th and 5th spectral bands.
- The matlab script for this example is multi_spectral_image.m.

Outline

Palette Images

- Palette images (indexed images) provide a simple way to reduce the amount of data needed to represent an image.
- The pixel values constitute a link to a lookup table of colors (also called a `color table`, `color map`, index register, `palette`).
- The lookup table contains as many entries as the range of possible values in the pixel, typically $8\ \mathrm{bits} = 256\ \mathrm{values}$.
- Each entry of the table maps the pixel value to the color, so there are three values, one for each of three color components.
- This approach would reduce data size to $1/3$ (plus size of the look up table).

Palette Images

- Palette images (indexed images) provide a simple way to reduce the amount of data needed to represent an image.
- The pixel values constitute a link to a lookup table of colors (also called a `color table`, `color map`, index register, `palette`).
- The lookup table contains as many entries as the range of possible values in the pixel, typically $8 \text{ bits} = 256 \text{ values}$.
- Each entry of the table maps the pixel value to the color, so there are three values, one for each of three color components.
- This approach would reduce data size to $1/3$ (plus size of the look up table).

Palette Images

- Palette images (indexed images) provide a simple way to reduce the amount of data needed to represent an image.
- The pixel values constitute a link to a lookup table of colors (also called a `color table`, `color map`, index register, `palette`).
- The lookup table contains as many entries as the range of possible values in the pixel, typically $8\,\text{bits} = 256\,\text{values}$.
- Each entry of the table maps the pixel value to the color, so there are three values, one for each of three color components.
- This approach would reduce data size to $1/3$ (plus size of the look up table).

Palette Images

- Palette images (indexed images) provide a simple way to reduce the amount of data needed to represent an image.
- The pixel values constitute a link to a lookup table of colors (also called a `color table`, `color map`, index register, `palette`).
- The lookup table contains as many entries as the range of possible values in the pixel, typically $8\,\text{bits} = 256\,\text{values}$.
- Each entry of the table maps the pixel value to the color, so there are three values, one for each of three color components.
- This approach would reduce data size to $1/3$ (plus size of the look up table).

Palette Images

- Palette images (indexed images) provide a simple way to reduce the amount of data needed to represent an image.
- The pixel values constitute a link to a lookup table of colors (also called a `color table`, `color map`, index register, `palette`).
- The lookup table contains as many entries as the range of possible values in the pixel, typically $8\,\text{bits} = 256\,\text{values}$.
- Each entry of the table maps the pixel value to the color, so there are three values, one for each of three color components.
- This approach would reduce data size to $1/3$ (plus size of the look up table).

Image Formats

- In the RGB color model, values for red, green and blue are provided.
- Image formats for raster images such as TIFF, PNG and GIF can store palette images.

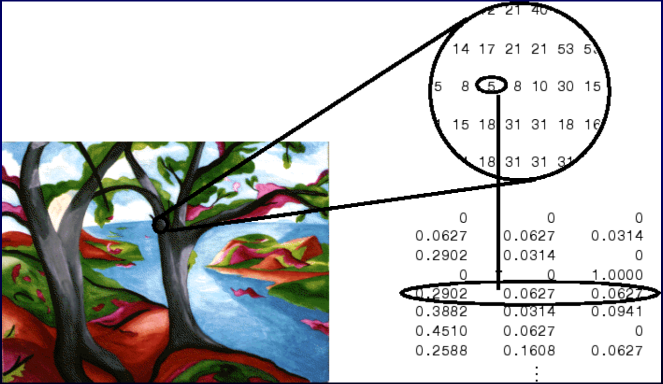

(image from www.mathworks.com)

Image Formats

- In the RGB color model, values for red, green and blue are provided.
- Image formats for raster images such as TIFF, PNG and GIF can store palette images.

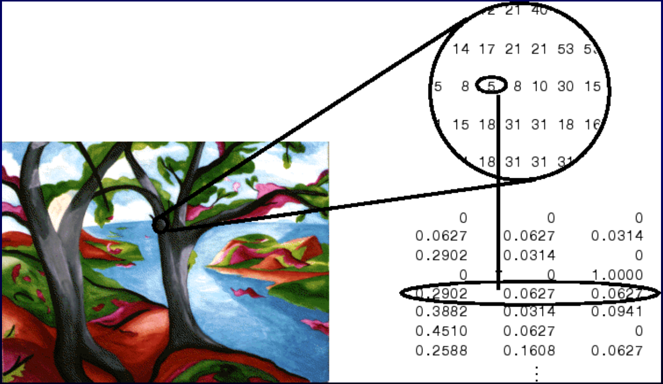

(image from www.mathworks.com)

Not Enough Colors?

- If the number of colors in the input image is less than or equal to the number of entries in the lookup table,
- then all colors can be selected and no loss of information occurs.
- Such images may be cartoon movies, or program outputs.
- In the more common case,
  - the number of colors in the image exceeds the number of entries in the lookup table
  - a subset of colors has to be chosen,
- and a loss of information occurs.

Not Enough Colors?

- If the number of colors in the input image is less than or equal to the number of entries in the lookup table,
- then all colors can be selected and no loss of information occurs.
- Such images may be cartoon movies, or program outputs.
- In the more common case,
  - the number of colors in the image exceeds the number of entries in the lookup table
  - a subset of colors has to be chosen,
- and a loss of information occurs.

Not Enough Colors?

- If the number of colors in the input image is less than or equal to the number of entries in the lookup table,
- then all colors can be selected and no loss of information occurs.
- Such images may be cartoon movies, or program outputs.
- In the more common case,
  - the number of colors in the image exceeds the number of entries in the lookup table
  - a subset of colors has to be chosen,
- and a loss of information occurs.

Not Enough Colors?

- If the number of colors in the input image is less than or equal to the number of entries in the lookup table,
- then all colors can be selected and no loss of information occurs.
- Such images may be cartoon movies, or program outputs.
- In the more common case,
  - the number of colors in the image exceeds the number of entries in the lookup table
  - a subset of colors has to be chosen,
- and a loss of information occurs.

Not Enough Colors?

- If the number of colors in the input image is less than or equal to the number of entries in the lookup table,
- then all colors can be selected and no loss of information occurs.
- Such images may be cartoon movies, or program outputs.
- In the more common case,
  - the number of colors in the image exceeds the number of entries in the lookup table
  - a subset of colors has to be chosen,
- and a loss of information occurs.

Not Enough Colors?

- If the number of colors in the input image is less than or equal to the number of entries in the lookup table,
- then all colors can be selected and no loss of information occurs.
- Such images may be cartoon movies, or program outputs.
- In the more common case,
  - the number of colors in the image exceeds the number of entries in the lookup table
  - a subset of colors has to be chosen,
- and a loss of information occurs.

Not Enough Colors?

- If the number of colors in the input image is less than or equal to the number of entries in the lookup table,
- then all colors can be selected and no loss of information occurs.
- Such images may be cartoon movies, or program outputs.
- In the more common case,
  - the number of colors in the image exceeds the number of entries in the lookup table
  - a subset of colors has to be chosen,
- and a loss of information occurs.

Color Selection

- This color selection may be done many ways.
- The simplest is to quantize color space regularly into cubes of the same size.
- In the 8 bit example,
  - there would be $8 \times 8 \times 8 = 256$ cubes.
  - for a green frog in green grass there will not be enough shades of green in the lookup table to display the image well.
- It is better to
  - check which colors appear in the image by creating histograms for all three color components
  - quantize them to provide more shades for colors which occur in the image frequently.
- If an image is converted to a palette representation then the `nearest` color in the lookup table is used to represent the color.

Color Selection

- This color selection may be done many ways.
- The simplest is to quantize color space regularly into cubes of the same size.
- In the 8 bit example,
  - there would be $8 \times 8 \times 8 = 256$ cubes.
  - for a green frog in green grass there will not be enough shades of green in the lookup table to display the image well.
- It is better to
  - check which colors appear in the image by creating histograms for all three color components
  - quantize them to provide more shades for colors which occur in the image frequently.
- If an image is converted to a palette representation then the `nearest` color in the lookup table is used to represent the color.

Color Selection

- This color selection may be done many ways.
- The simplest is to quantize color space regularly into cubes of the same size.
- In the 8 bit example,
  - there would be $8 \times 8 \times 8 = 256$ cubes.
  - for a green frog in green grass there will not be enough shades of green in the lookup table to display the image well.
- It is better to
  - check which colors appear in the image by creating histograms for all three color components
  - quantize them to provide more shades for colors which occur in the image frequently.
- If an image is converted to a palette representation then the `nearest` color in the lookup table is used to represent the color.

Color Selection

- This color selection may be done many ways.
- The simplest is to quantize color space regularly into cubes of the same size.
- In the 8 bit example,
  - there would be $8 \times 8 \times 8 = 256$ cubes.
  - for a green frog in green grass there will not be enough shades of green in the lookup table to display the image well.
- It is better to
  - check which colors appear in the image by creating histograms for all three color components
  - quantize them to provide more shades for colors which occur in the image frequently.
- If an image is converted to a palette representation then the `nearest` color in the lookup table is used to represent the color.

Color Selection

- This color selection may be done many ways.
- The simplest is to quantize color space regularly into cubes of the same size.
- In the 8 bit example,
  - there would be $8 \times 8 \times 8 = 256$ cubes.
  - for a green frog in green grass there will not be enough shades of green in the lookup table to display the image well.
- It is better to
  - check which colors appear in the image by creating histograms for all three color components
  - quantize them to provide more shades for colors which occur in the image frequently.
- If an image is converted to a palette representation then the `nearest` color in the lookup table is used to represent the color.

Color Selection

- This color selection may be done many ways.
- The simplest is to quantize color space regularly into cubes of the same size.
- In the 8 bit example,
  - there would be $8 \times 8 \times 8 = 256$ cubes.
  - for a green frog in green grass there will not be enough shades of green in the lookup table to display the image well.
- It is better to
  - check which colors appear in the image by creating histograms for all three color components
  - quantize them to provide more shades for colors which occur in the image frequently.
- If an image is converted to a palette representation then the `nearest` color in the lookup table is used to represent the color.

Color Selection

- This color selection may be done many ways.
- The simplest is to quantize color space regularly into cubes of the same size.
- In the 8 bit example,
  - there would be $8 \times 8 \times 8 = 256$ cubes.
  - for a green frog in green grass there will not be enough shades of green in the lookup table to display the image well.
- It is better to
  - check which colors appear in the image by creating histograms for all three color components
  - quantize them to provide more shades for colors which occur in the image frequently.
- If an image is converted to a palette representation then the `nearest` color in the lookup table is used to represent the color.

Color Selection

- This color selection may be done many ways.
- The simplest is to quantize color space regularly into cubes of the same size.
- In the 8 bit example,
  - there would be $8 \times 8 \times 8 = 256$ cubes.
  - for a green frog in green grass there will not be enough shades of green in the lookup table to display the image well.
- It is better to
  - check which colors appear in the image by creating histograms for all three color components
  - quantize them to provide more shades for colors which occur in the image frequently.
- If an image is converted to a palette representation then the `nearest` color in the lookup table is used to represent the color.

Color Selection

- This color selection may be done many ways.
- The simplest is to quantize color space regularly into cubes of the same size.
- In the 8 bit example,
  - there would be $8 \times 8 \times 8 = 256$ cubes.
  - for a green frog in green grass there will not be enough shades of green in the lookup table to display the image well.
- It is better to
  - check which colors appear in the image by creating histograms for all three color components
  - quantize them to provide more shades for colors which occur in the image frequently.
- If an image is converted to a palette representation then the `nearest` color in the lookup table is used to represent the color.

Pseudocolor

- The term `pseudocolor` is usually used when an original image is gray-level and is displayed in color;
  - this is often done to exploit the color discriminatory power of human vision.
- The same palette machinery as described above is used for this purpose;
  - a palette is loaded into the lookup table which visualizes the particular gray-scale image the best.
- It could either enhance local changes, or might provide various views of the image.
- Which palette to choose is an issue.

Pseudocolor

- The term `pseudocolor` is usually used when an original image is gray-level and is displayed in color;
  - this is often done to exploit the color discriminatory power of human vision.
- The same palette machinery as described above is used for this purpose;
  - a palette is loaded into the lookup table which visualizes the particular gray-scale image the best.
- It could either enhance local changes, or might provide various views of the image.
- Which palette to choose is an issue.

Pseudocolor

- The term `pseudocolor` is usually used when an original image is gray-level and is displayed in color;
  - this is often done to exploit the color discriminatory power of human vision.
- The same palette machinery as described above is used for this purpose;
  - a palette is loaded into the lookup table which visualizes the particular gray-scale image the best.
- It could either enhance local changes, or might provide various views of the image.
- Which palette to choose is an issue.

Pseudocolor

- The term `pseudocolor` is usually used when an original image is gray-level and is displayed in color;
  - this is often done to exploit the color discriminatory power of human vision.
- The same palette machinery as described above is used for this purpose;
  - a palette is loaded into the lookup table which visualizes the particular gray-scale image the best.
- It could either enhance local changes, or might provide various views of the image.
- Which palette to choose is an issue.

Pseudocolor

- The term `pseudocolor` is usually used when an original image is gray-level and is displayed in color;
  - this is often done to exploit the color discriminatory power of human vision.
- The same palette machinery as described above is used for this purpose;
  - a palette is loaded into the lookup table which visualizes the particular gray-scale image the best.
- It could either enhance local changes, or might provide various views of the image.
- Which palette to choose is an issue.

Pseudocolor

- The term `pseudocolor` is usually used when an original image is gray-level and is displayed in color;
  - this is often done to exploit the color discriminatory power of human vision.
- The same palette machinery as described above is used for this purpose;
  - a palette is loaded into the lookup table which visualizes the particular gray-scale image the best.
- It could either enhance local changes, or might provide various views of the image.
- Which palette to choose is an issue.

Outline

Color Constancy: a challenging problem

- Consider the situation in which the same surface is seen under different illumination.
- The same surface colors are shown fully illuminated and in shadow.
- The human vision system is able to abstract to a certain degree from the illumination changes and perceive several instances of a particular color as the same.
- This phenomenon is called `color constancy`.

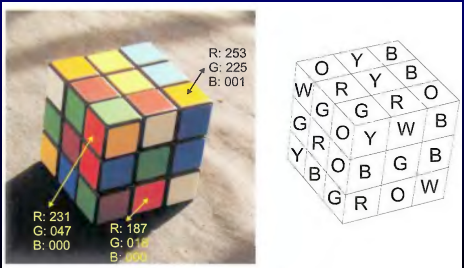

Color Constancy: a challenging problem

- Consider the situation in which the same surface is seen under different illumination.
- The same surface colors are shown fully illuminated and in shadow.
- The human vision system is able to abstract to a certain degree from the illumination changes and perceive several instances of a particular color as the same.
- This phenomenon is called `color constancy`.

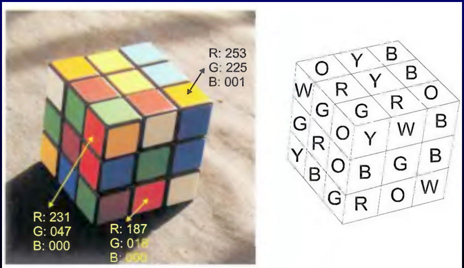

Color Constancy: a challenging problem

- Consider the situation in which the same surface is seen under different illumination.
- The same surface colors are shown fully illuminated and in shadow.
- The human vision system is able to abstract to a certain degree from the illumination changes and perceive several instances of a particular color as the same.
- This phenomenon is called `color constancy`.

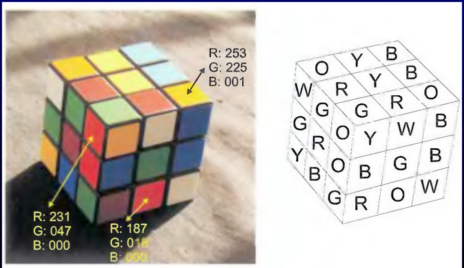

Color Constancy: a challenging problem

- Consider the situation in which the same surface is seen under different illumination.
- The same surface colors are shown fully illuminated and in shadow.
- The human vision system is able to abstract to a certain degree from the illumination changes and perceive several instances of a particular color as the same.
- This phenomenon is called `color constancy`.

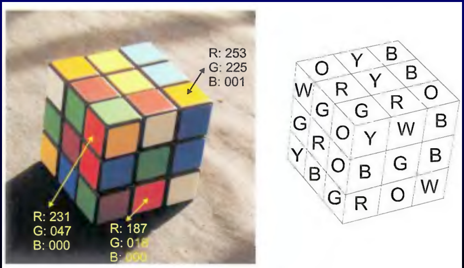

References I

- D. R. Wilson and T. R. Martinez (1997).
  Improved heterogeneous distance functions.
  Journal of Artificial Intelligence Research.
- Frisby, J. P. (1979).
  Seeing — Illusion, Brain and Mind.
  Oxford University Press, Oxford.
- Horn, B. R. (1986).
  Robot Vision.
  M.I.T. Press, Cambridge.
- Knuth, D. E. (1981).
  Seminumerical Algorithms.
  Addison-Wesley, Reading, Massachusetts, 2nd edition.
  volume 2 of The art of Computer Programming.

References II

- Kovalevski, V. A. (1989).
  Finite topology as applied to image analysis.
  Computer Vision, Graphics and Image Processing, 46:141 – 161.
- Kutter, M. (1999).
  Digital Image Watermarking: hiding information in images.
  ph.d dissertation, Départment d'Électricité, Ècole Polytechnique Fédérale de Lausanne.
- Oppenheim, A. V., Willsky, A. S., and Nawwab, S. H. (1997).
  Signals and Systems.
  Prentice Hall, Inc.

References III

- Pavlidis, T. (1977).
  Structural Pattern Recognition.
  Springer Verlag, Berlin, Heidelberg, New York.
- Pratt, W. K. (1978).
  Digital Image Processing.
  John Wiley & Sons, Inc, New York.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1992).
  Numerical Recipes in C: The Art of Scientifc Computing.
  Cambridge University Press, 2nd edition.
- Shannon, C. E. (1948).
  A mathematical theory of communication.
  The Bell System Technical Journal, 27:379 – 656 and 623 – 656.

References IV

- Wandell, B. (1995).
  Foundation of Vision.
  Sinauer Associates.
- Wang, Z., Bovik, A. C., Sheikh, H. R., Member, S., Simoncelli, E. P., and Member, S. (2004).
  Image quality assessment: From error visibility to structural similarity.
  IEEE Transactions on Image Processing, 13:600–612.
- Wang, Z. J. and Wu, S. D. (1991).
  Imaging Optics.
  Science Press, Beijing.
  in Chinese.