Technical University of Lodz
Institute of Electronics

Piotr M. Szczypiński

Image Processing and Computer Graphics
ENHANCEMENT part 2

Przetwarzanie obrazów i grafika komputerowa
POPRAWA JAKOŚCI cz. 2

Median filter

The median ***m*** of a set of values (e.g. image pixels in the filtering mask) is such that half the elements in the set are less than ***m*** and other half are grater than ***m***.

$$x(n)=\{1, 5, -7, 101, -25, 3, 0, 11, 7\}$$

Sorted sequence of elements:

$$x_s(n)=\{-25, -7, 0, 1, \mathbf{3}, 5, 7, 11, 101\}$$

Median filter

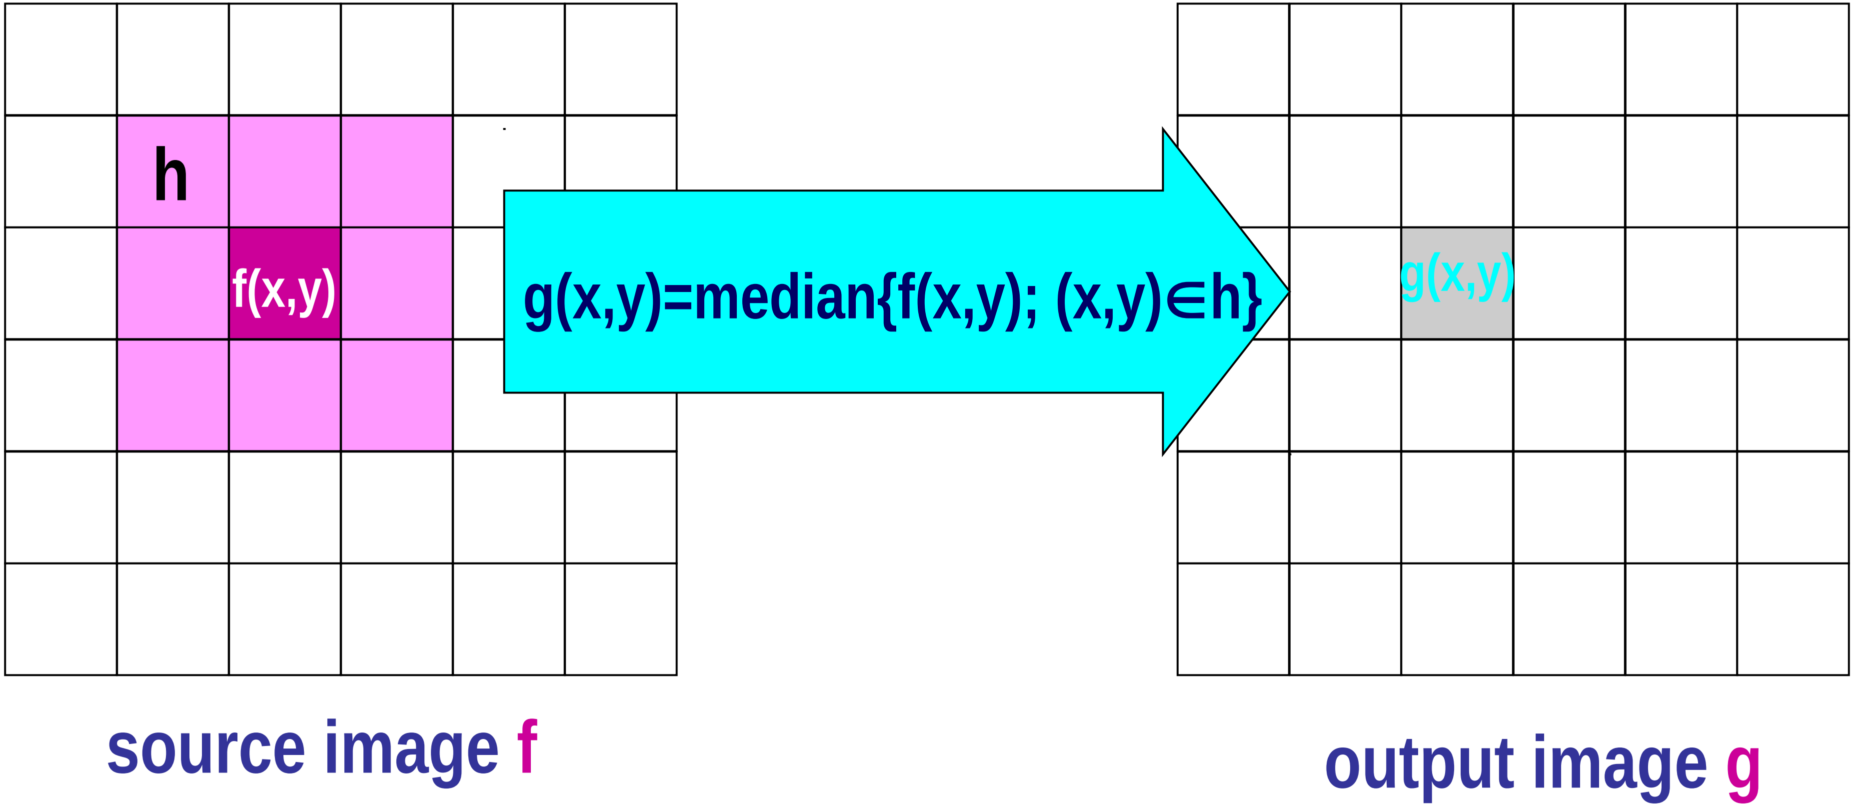

Median filter

From Wikipedia, the free encyclopedia

In image processing it is usually necessary to perform a high degree of noise reduction in an image before performing higher-level processing steps, such as edge detection. The **median filter** is a non-linear digital filtering technique, often used to remove noise from images or other signals. The idea is to examine a sample of the input and decide if it is representative of the signal. This is performed using a window consisting of an odd number of samples. The values in the window are sorted into numerical order; the median value, the sample in the center of the window, is selected as the output. The oldest sample is discarded, a new sample acquired, and the calculation repeats.

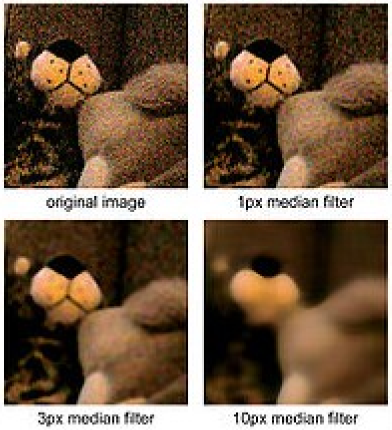

Example of 3 median filters of varying radii applied to the same noisy photograph. Implemented in Adobe Photoshop.

Median filtering is a common step in image processing. It is particularly useful to reduce speckle noise and salt and pepper noise. Its edge-preserving nature makes it useful in cases where edge blurring is undesirable.

Contents [hide]

navigation
- Main page
- Contents
- Featured content
- Current events
- Random article

search

Go | Search

interaction
- About Wikipedia
- Community portal
- Recent changes
- Contact Wikipedia
- Donate to Wikipedia
- Help

Median filter

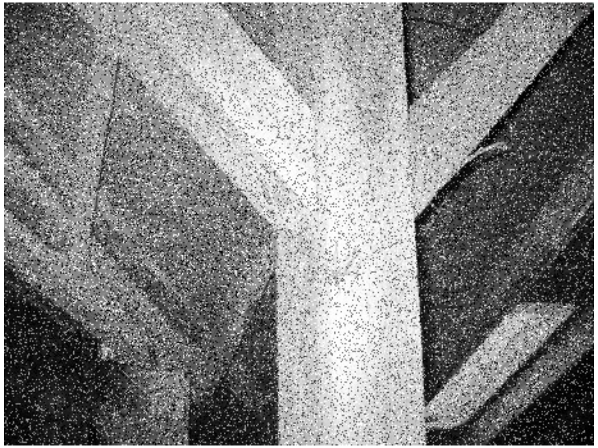

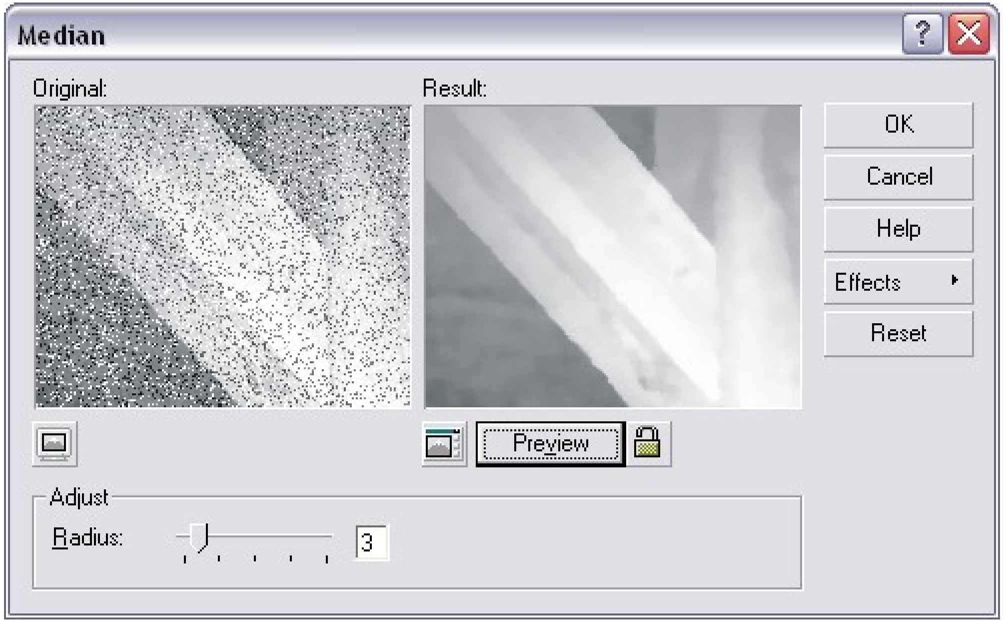

Image morphology

Median filter…

What about minimum or maximum filters?

Image morphology - erosion

Minimum filter – morphological erosion

Maximum filter – morphological dilation

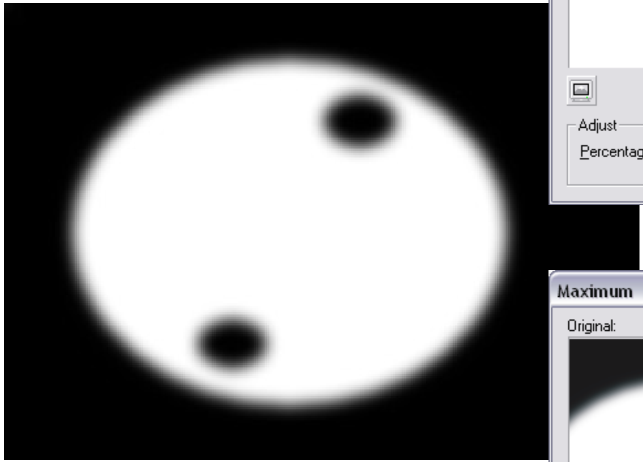

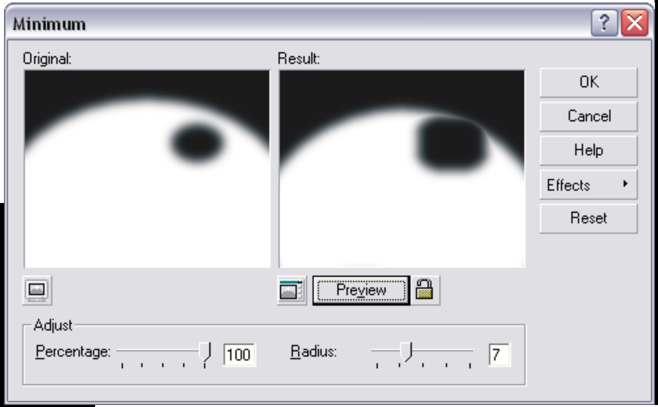

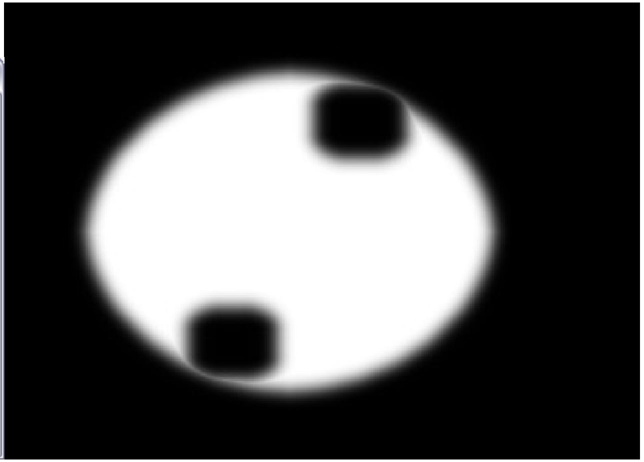

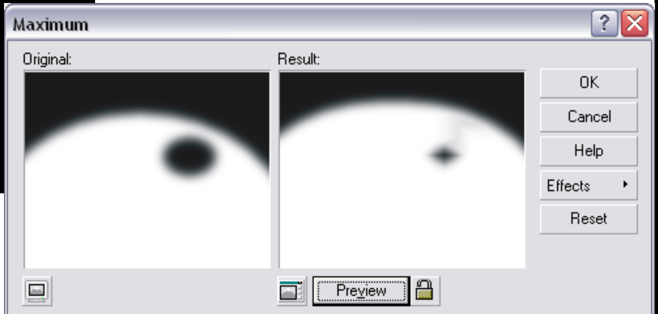

Image morphology - erosion

Minimum filter – morphological erosion

Maximum filter – morphological dilation

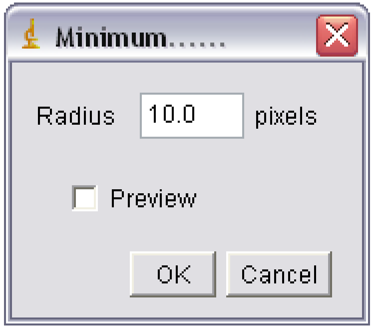

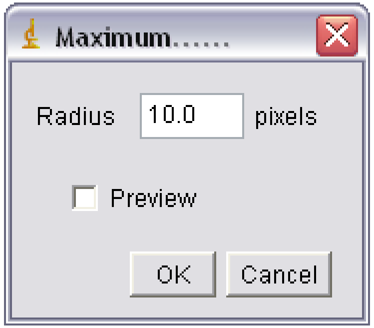

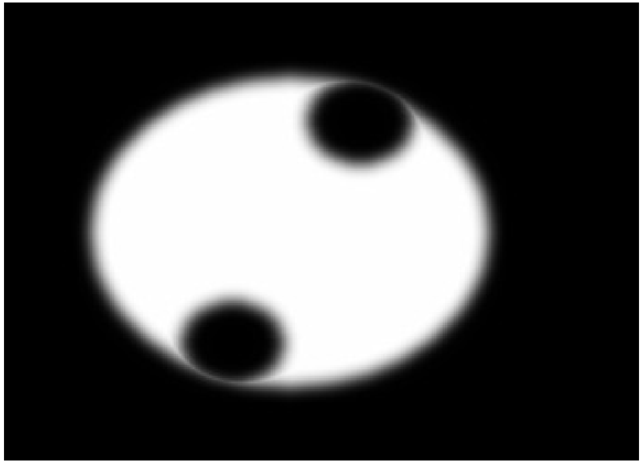

Mathematical morphology

From Wikipedia, the free encyclopedia

For other uses, see Morphology.

Mathematical morphology (MM) is a theory and technique for the analysis and processing of geometrical structures, based on set theory, lattice theory, topology, and random functions. MM is most commonly applied to digital images, but it can be employed as well on graphs, surface meshes, solids, and many other spatial structures.

Topological and geometrical continuous-space concepts such as size, shape, convexity, connectivity, and geodesic distance, can be characterized by MM on both continuous and discrete spaces. MM is also the foundation of morphological image processing, which consists of a set of operators that transform images according to the above characterizations.

MM was originally developed for binary images, and was later extended to grayscale functions and images. The subsequent generalization to complete lattices is widely accepted today as MM's theoretical foundation.

A shape (in blue) and its morphological dilation (in green) and erosion (in yellow) by a diamond-shape structuring element.

Contents [hide]

Erosion (morphology)

From Wikipedia, the free encyclopedia

For use of "Erosion" in dermatopathology, see Erosion_(dermatopathology)

Erosion is one of two fundamental operations (the other being dilation) in Morphological image processing from which all other morphological operations are based. It was originally defined for binary images, later being extended to grayscale images, and subsequently to complete lattices.

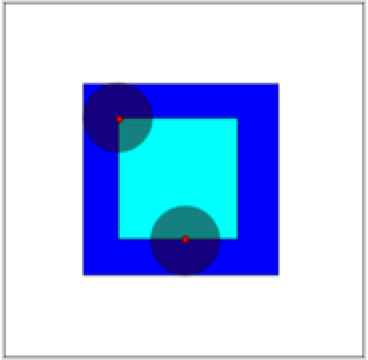

The erosion of the dark-blue square by a disk, resulting in the light-blue square.

Contents [hide]

Binary erosion [edit]

In binary morphology, an image is viewed as a subset of an Euclidean space $\mathbb{R}^d$ or the integer grid $\mathbb{Z}^d$, for some dimension *d*.

The basic idea in binary morphology is to probe an image with a simple, pre-defined shape, drawing conclusions on how this shape fits

Dilation (morphology)

From Wikipedia, the free encyclopedia

Dilation is one of the basic operations in mathematical morphology. Originally developed for binary images, it has been expanded first to grayscale images, and then to complete lattices. The dilation operation usually uses a structuring element for probing and expanding the shapes contained in the input image.

Contents [hide]

Binary Operator

[edit]

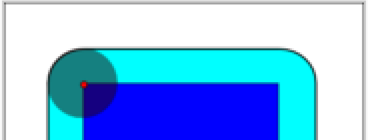

In binary morphology, dilation is a shift-invariant (translation invariant) operator, strongly related to the Minkowski addition.

A binary image is viewed in mathematical morphology as a subset of an Euclidean space R^d or the integer grid Z^d, for some dimension d. Let E be an Euclidean space or an integer grid, A a binary

Image morphology - erosion

Opening = erosion + dilation

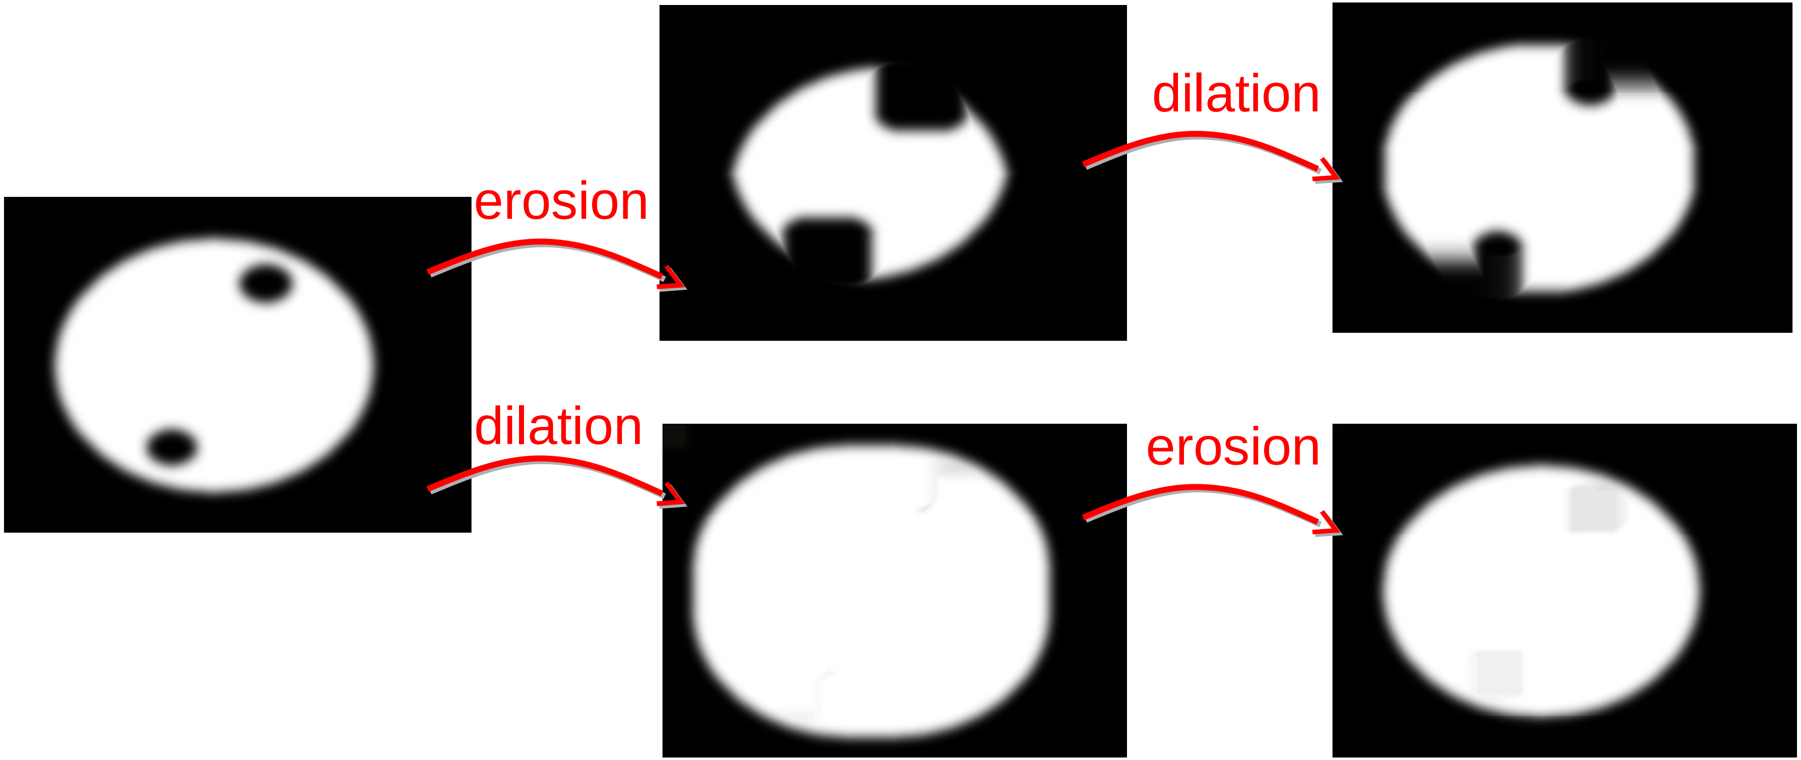

Closing = dilation + erosion

Image morphology - erosion

Opening = erosion + dilation

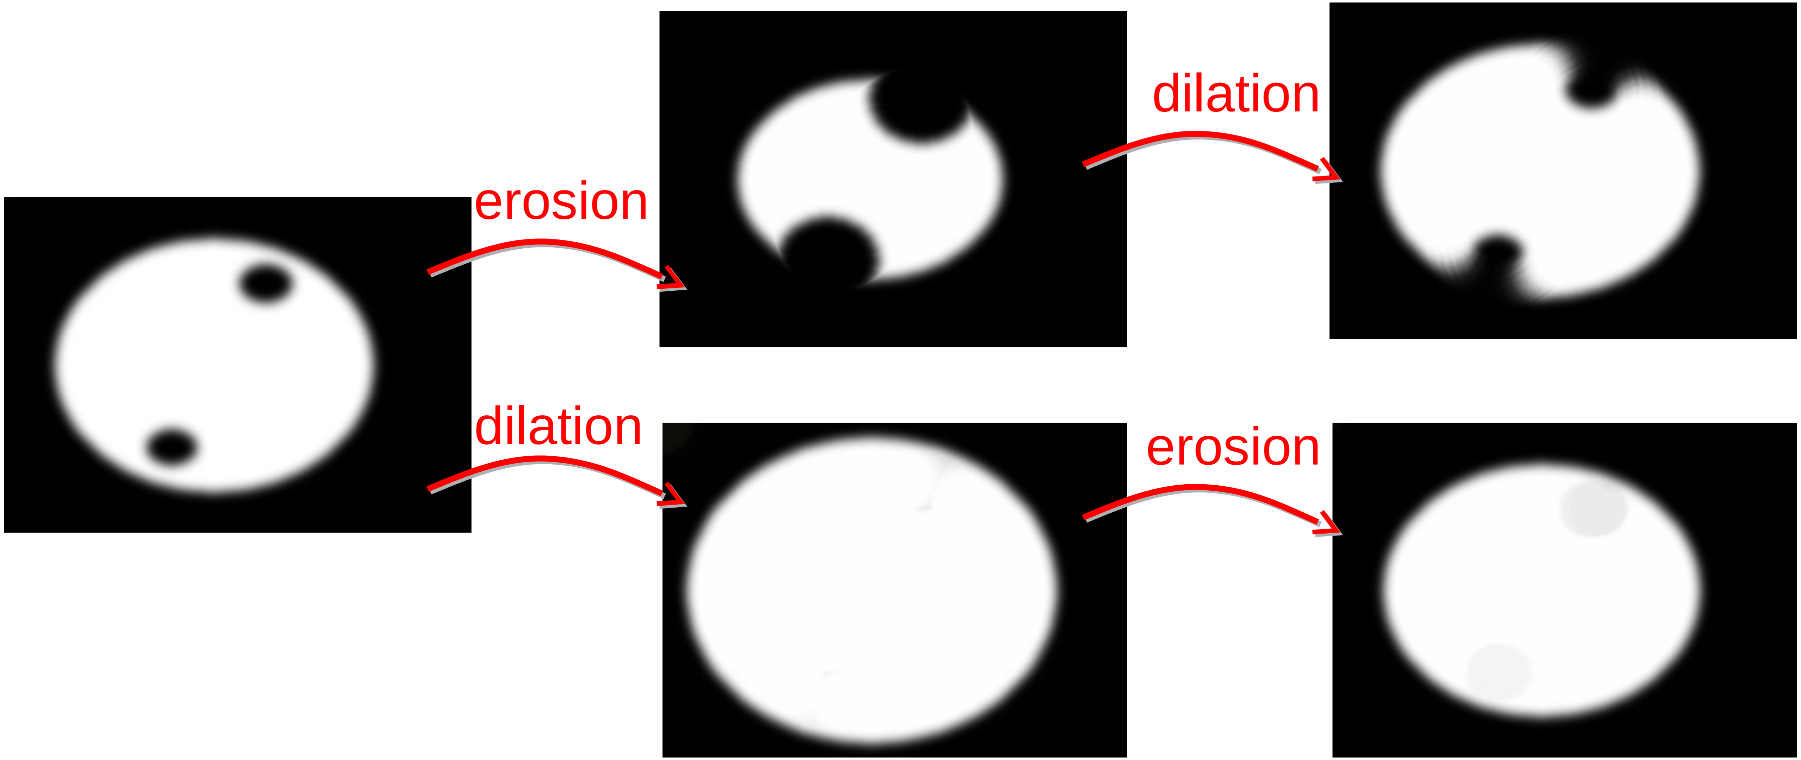

Closing = dilation + erosion

Closing (morphology)

From Wikipedia, the free encyclopedia

In mathematical morphology, the **closing** of a set (binary image) *A* by a structuring element *B* is the erosion of the dilation of that set,

$$A \bullet B = (A \oplus B) \ominus B,$$

where $\oplus$ and $\ominus$ denote the dilation and erosion, respectively.

In image processing, closing is, together with opening, the basic workhorse of morphological noise removal. Opening removes small objects, while closing removes small holes.

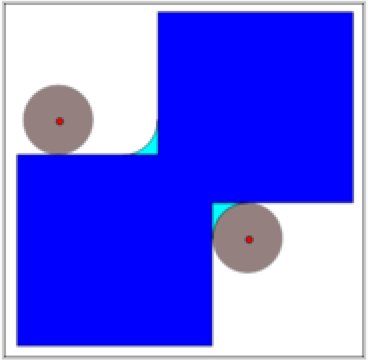

The closing of the dark-blue shape (union of two squares) by a disk, resulting in the union of the dark-blue shape and the light-blue areas.

Properties [edit]

- It is idempotent, that is, $(A \bullet B) \bullet B = A \bullet B$.
- It is increasing, that is, if $A \subseteq C$, then $A \bullet B \subseteq C \bullet B$.
- It is extensive, i.e., $A \subseteq A \bullet B$.
- It is translation invariant.

Bibliography [edit]

- *Image Analysis and Mathematical Morphology* by Jean Serra, ISBN 0126372403 (1982)
- *Image Analysis and Mathematical Morphology, Volume 2: Theoretical Advances* by Jean Serra, ISBN 0-12-637241-1 (1988)

Opening (morphology)

From Wikipedia, the free encyclopedia

In mathematical morphology, **opening** is the dilation of the erosion of a set A by a structuring element B:

$$A \circ B = (A \ominus B) \oplus B,$$

where $\ominus$ and $\oplus$ denote erosion and dilation, respectively.

Together with closing, the opening serves in computer vision and image processing as a basic workhorse of morphological noise removal. Opening removes small objects from the foreground (usually taken as the dark pixels) of an image, placing them in the background, while closing removes small holes in the foreground, changing small islands of background into foreground. These techniques can also be used to find specific shapes in an image. Opening can be used to find things into which a specific structuring element can fit (edges, corners, ...).

One can think of *B* sweeping around the inside of the boundary of *A*, so that it does not extend beyond the boundary, and shaping the *A* boundary around the boundary of the element.

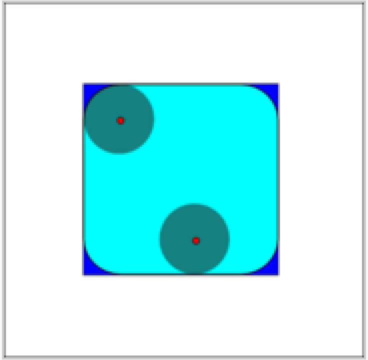

The opening of the dark-blue square by a disk, resulting in the light-blue square with round corners.

Properties

[edit]

- Opening is idempotent, that is, $(A \circ B) \circ B = A \circ B$.
- Opening is increasing, that is, if $A \subseteq C$, then $A \circ B \subseteq C \circ B$.

Image resampling (zooming)

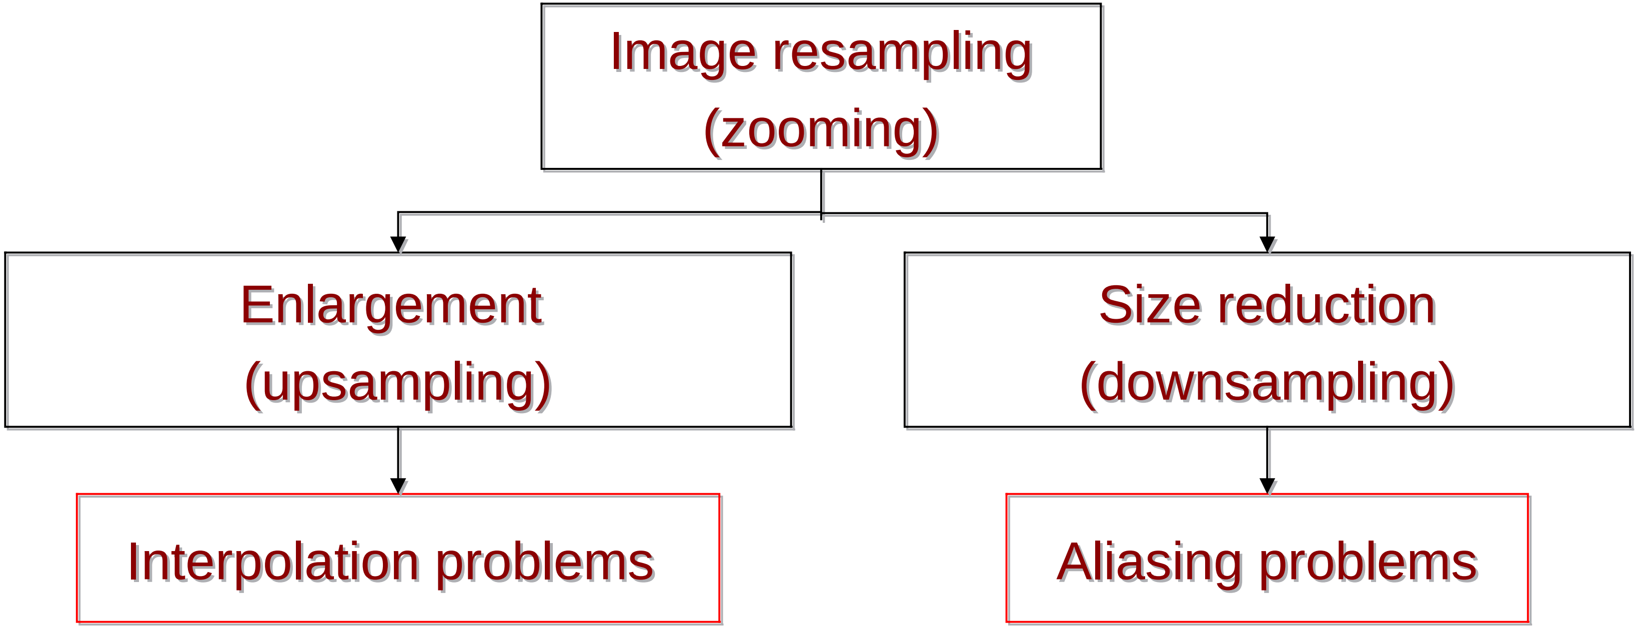

Resampling

From Wikipedia, the free encyclopedia

For resampling methods in statistics, see Resampling (statistics).

It has been suggested that this article be split into articles entitled *Resampling (audio)* and *Resampling (bitmap)*, accessible from a disambiguation page. (Discuss)

Resampling is the digital process of changing the sample rate or dimensions of digital imagery or audio by temporally or areally analysing and sampling the original data.

Contents [hide]

Audio [edit]

Audio resampling is also called sample rate conversion. This operation in digital signal processing involves converting a sampled signal from one sampling frequency to another. For instance, the output waveform of a digital audio workstation that was processed at 96 kHz must be resampled to 44.1 kHz to be placed on a Compact Disc. The article **Sample rate conversion** explains how this is done.

Downsampling

Just take every n-th pixel!

SIMPLE!?

Downsampling

From Wikipedia, the free encyclopedia

For the computer graphics technique, see Subsampling.

In signal processing, **downsampling** (or "subsampling") is the process of reducing the sampling rate of a signal. This is usually done to reduce the data rate or the size of the data.

The downsampling factor (commonly denoted by *M*) is usually an integer or a rational fraction greater than unity. This factor multiplies the sampling time or, equivalently, divides the sampling rate. For example, if compact disc audio is downsampled by a factor of 5/4 then the resulting sampling rate goes from 44,100 Hz to 35,280 Hz, which reduces the bit rate from 1,411,200 bit/s to 1,128,960 bit/s (assuming that a 32 bit value is sampled each time).

Contents [hide]

Maintaining the sampling theorem criterion

[edit]

Since downsampling reduces the sampling rate, we must be careful to make sure the Shannon-Nyquist sampling theorem criterion is maintained. If the sampling theorem is not satisfied then the resulting digital signal will have aliasing. To ensure that the sampling

Nyquist frequency - Wikipedia, the free encyclopedia - Windows Internet Explorer

http://en.wikipedia.org/wiki/Nyquist_frequency

Plik Edycja Widok Ulubione Narzędzia Pomoc

Resampling - ... Downsampling... Upsampling - ... Nyquist fr... Nyquist–Shan... Strona Narzędzia

Log in / create account

article discussion edit this page history

Nyquist frequency

From Wikipedia, the free encyclopedia

Not to be confused with Nyquist rate.

The **Nyquist frequency**, named after the Swedish-American engineer Harry Nyquist or the Nyquist–Shannon sampling theorem, is half the sampling frequency of a discrete signal processing system.[1][2] It is sometimes called the **folding frequency**, or the **cut-off frequency** of a sampling system.[3]

The sampling theorem shows that aliasing can be avoided if the Nyquist frequency is greater than the bandwidth, or maximum component frequency, of the signal being sampled.

The Nyquist frequency should not be confused with the *Nyquist rate*, which is the lower bound of the sampling frequency that satisfies the Nyquist sampling criterion for a given signal or family of signals. This lower bound is twice the bandwidth or maximum component frequency of the signal. *Nyquist rate*, as commonly used with respect to sampling, is a property of a continuous-time signal, not of a system, whereas *Nyquist frequency* is a property of a discrete-time system, not of a signal. The domain of the signals does not have to be time, though that is common, leading to Nyquist frequency in hertz; for example, an image sampling system has a Nyquist frequency expressed in units such as cycles per meter.

Contents [hide]

Nyquist-Shannon sampling theorem - Wikipedia, the free encyclopedia - Windows Internet Explorer

http://en.wikipedia.org/wiki/Nyquist%E2%80%93Shannon_sampling_theorem

Plik Edycja Widok Ulubione Narzędzia Pomoc

Resampling - … | Downsampling… | Upsampling - … | Nyquist frequ… | Nyquist–S…

Strona Narzędzia

Log in / create account

article | discussion | edit this page | history

WIKIPEDIA
The Free Encyclopedia

navigation

- Main page
- Contents
- Featured content
- Current events
- Random article

search

Go | Search

interaction

- About Wikipedia
- Community portal
- Recent changes
- Contact Wikipedia
- Donate to Wikipedia
- Help

Nyquist–Shannon sampling theorem

From Wikipedia, the free encyclopedia

The **Nyquist–Shannon sampling theorem** is a fundamental result in the field of information theory, in particular telecommunications and signal processing. Sampling is the process of converting a signal (for example, a function of continuous time or space) into a numeric sequence (a function of discrete time or space). The theorem states:[1]

> If a function $x(t)$ contains no frequencies higher than B cps, it is completely determined by giving its ordinates at a series of points spaced $1/(2B)$ seconds apart.

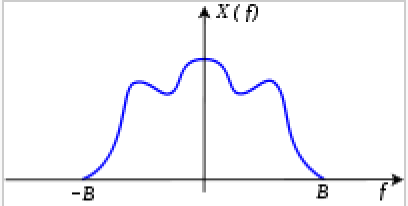

Fig.1: Hypothetical spectrum of a **bandlimited signal** as a function of frequency

In essence the theorem shows that an analog signal that has been sampled can be perfectly reconstructed from the samples if the sampling rate exceeds $2B$ samples per second, where B is the highest frequency in the original signal. If a signal contains a component at exactly B hertz, then samples spaced at exactly $1/(2B)$ seconds do not completely determine the signal, Shannon's statement notwithstanding.

More recent statements of the theorem are sometimes careful to exclude the equality condition; that is, the condition is if $x(t)$ contains no frequencies higher than *or equal to* B; this condition is equivalent to Shannon's except when the function includes a steady sinusoidal component at exactly frequency B.

The assumptions necessary to prove the theorem form a mathematical model that is only an idealization of any real-world situation. The conclusion that perfect reconstruction is possible is mathematically correct for the model but only an approximation for actual signals and actual sampling techniques.

The theorem also leads to a formula for reconstruction of the original signal. The constructive proof of the theorem leads to an

Aliasing

[edit]

Main article: Aliasing

The Poisson summation formula indicates that the samples of function **x(t)** are sufficient to create a *periodic extension* of function **X(f)**. The result is:

$$X_s(f) \stackrel{\text{def}}{=} \sum_{k=-\infty}^{\infty} X\left(f - k f_s\right) = T \sum_{n=-\infty}^{\infty} x(nT)\ e^{-i 2\pi n T f}. \qquad \textbf{(Eq.1)}$$

As depicted in Figures 3, 4, and 8, copies of **X(f)** are shifted by multiples of f_s and combined by addition.

If the sampling condition is not satisfied, adjacent copies overlap, and it is not possible in general to discern an unambiguous **X(f)**. Any frequency component above $f_s/2$ is indistinguishable from a lower-frequency component, called an *alias*, associated with one of the copies. The reconstruction technique described below produces the alias, rather than the original component, in such cases.

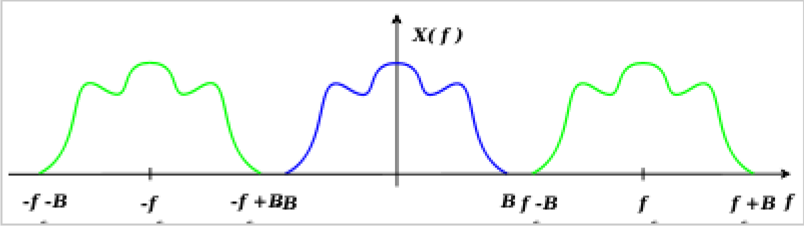

Fig.3: Hypothetical spectrum of a properly sampled bandlimited signal (blue) and images (green) that do not overlap. A "*brick-wall*" low-pass filter can remove the images and leave the original spectrum, thus recovering the original signal from the samples.

For a sinusoidal component of exactly half the sampling frequency, the component will in general alias to another sinusoid of the same frequency, but with a different phase

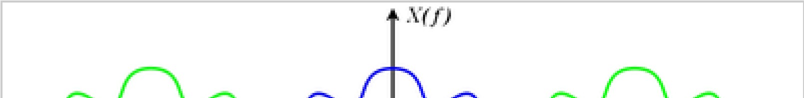

Application to multivariable signals and images [edit]

Hypothetical spectrum of a marginally sufficiently sampled bandlimited signal (blue), $X_A(f)$, where the images (green) narrowly do not overlap. But the overall sampled spectrum of $X_A(f)$ is identical to the overall inadequately sampled spectrum of $X(f)$ (top) because the sum of baseband and images are the same in both cases. The discrete sampled signals $x_A[n]$ and $x[n]$ are also identical. It is not possible, just from examining the spectra (or the sampled signals), to tell the two situations apart. If this were an audio signal, $x_A[n]$ and $x[n]$ would sound the same and the presumed "*properly*" sampled $x_A[n]$ would be the *alias* of $x[n]$ since the ... he spectrum $X(f)$.

Fig.5: Subsampled image showing a Moiré pattern

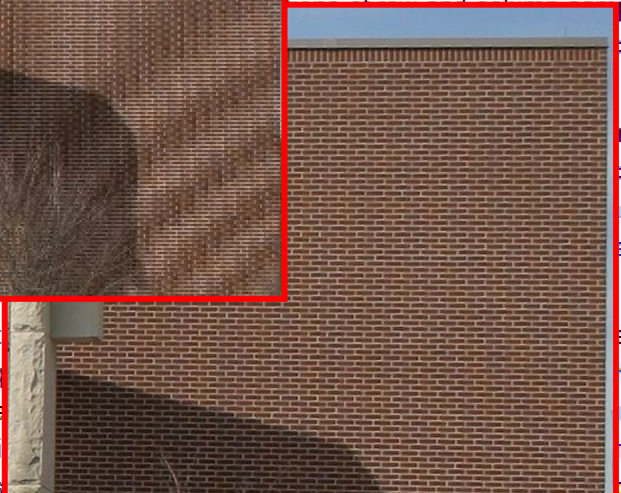

The sampling theorem is usually formulated for functions of a single variable. Consequently, the theorem is directly applicable to ti... dependent signals and is normally formulated in th... straightforward way to fu... represented as two-dime... intensities of pixels (pictu... locations. As a result, im... uniquely — one for the ro... eorem can be extended in a ... Grayscale images, for example, are often ... umbers representing the relative ... ple ... h pixel

Color images typically c... each of the three primary... vectors for colors include... black (CMYK) may repre... functions over a two-dime... represent ... es using 3- ... , yellow, and ... alued

Similar to one-dimensional discrete-t... resolution, or pixel density, is inadeq... frequencies (in other words, the dista... when it is sampled by the camera's i... "solution" to higher sampling in the s... e sampling ... hirt with high ... of the shirt ... . The ... the shirt

Upsampling

This must be simple!

Let’s multiply pixels!

Upsampling

From Wikipedia, the free encyclopedia

Upsampling is the process of increasing the sampling rate of a signal. For instance, upsampling raster images such as photographs means increasing the resolution of the image.

The upsampling factor (commonly denoted by *L*) is usually an integer or a rational fraction greater than unity. This factor multiplies the sampling rate or, equivalently, divides the sampling period. For example, if compact disc audio is upsampled by a factor of 5/4 then the resulting sampling rate goes from 44,100 Hz to 55,125 Hz.

Contents [hide]

Sampling theorem satisfaction [edit]

The upsampled signal satisfies the Nyquist–Shannon sampling theorem if the original signal does.

For an aesthetically pleasing upsample, an interpolation filter is required; in both upsampling and downsampling, such a low-pass filter implements anti-aliasing.

Nearest neighbor pixel replication

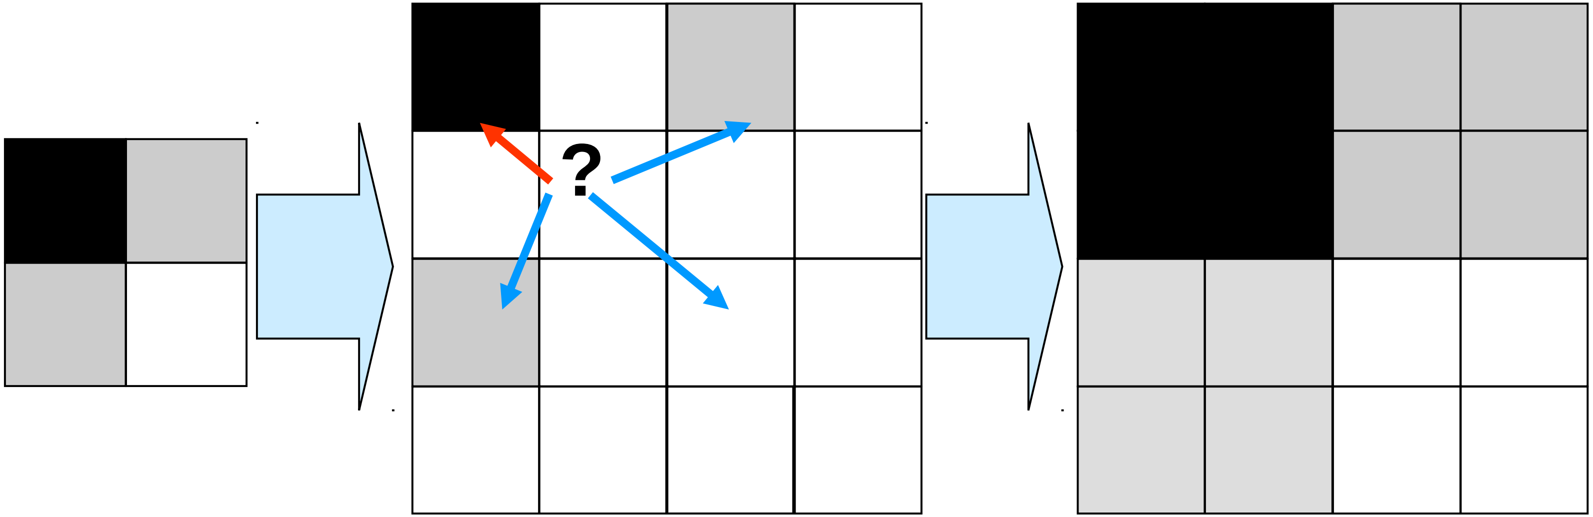

Zero-order interpolation

Nearest-neighbor interpolation

From Wikipedia, the free encyclopedia

Nearest-neighbor interpolation (also known as **proximal interpolation** or **point sampling** in some contexts) is a simple method of multivariate interpolation in 1 or more dimensions. Interpolation is the problem of approximating the value for a non-given point in some space, when given some values of points *around* that point. The nearest neighbor algorithm simply selects the value of the nearest point, and does not consider the values of other neighboring points at all, yielding a piecewise-constant interpolant. The algorithm is very simple to implement, and is commonly used (usually along with mipmapping) in real-time 3D rendering to select color values for a textured surface.

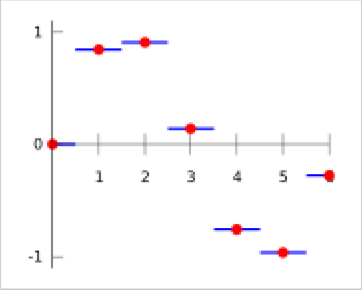

Nearest neighbor interpolation (blue lines) in one dimension on a dataset (red points).

Connection to Voronoi diagram [edit]

For a given set of points in space, a Voronoi diagram is a decomposition of space into cells, one for each given point, so that anywhere in space, the closest given point is inside the cell. This is equivalent to nearest neighbour interpolation, by assigning the function value at the given point to all the points inside the cell. The figures on the right side show by colour the shape of the cells.

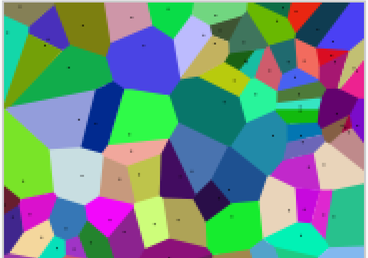

See also [edit]

- Interpolation
- Nearest neighbor search

navigation
- Main page
- Contents
- Featured content
- Current events
- Random article

search
Go Search

interaction
- About Wikipedia
- Community portal
- Recent changes
- Contact Wikipedia
- Donate to Wikipedia
- Help

Nearest neighbor pixel replication

Pixelation! ☹

Bilinear interpolation

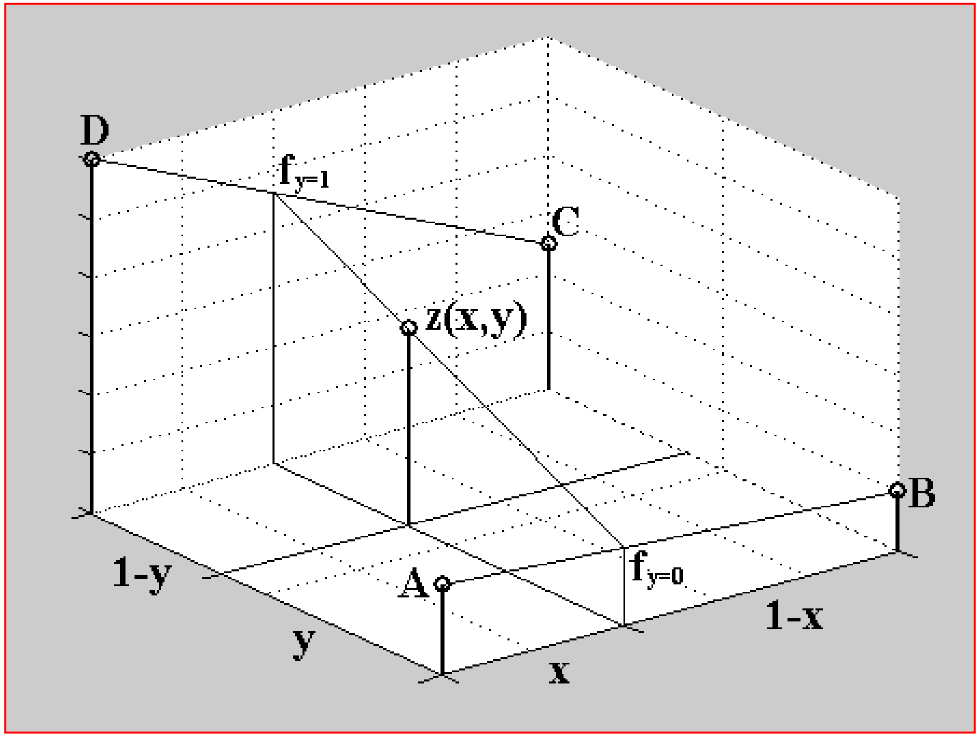

$x \in <0,1>$

$y \in <0,1>$

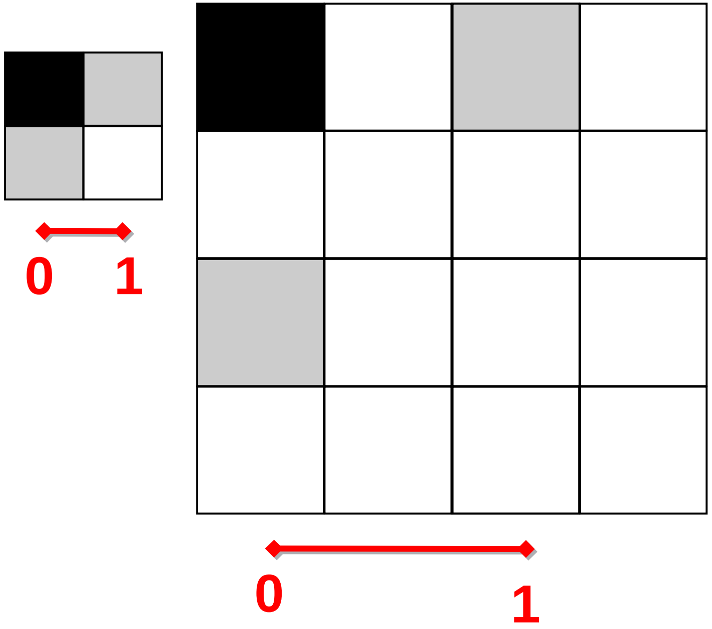

$$z(x,y) = A(1-x)(1-y) + B(x)(1-y) + C(x)(y) + D(1-x)(y)$$

Bilinear interpolation

From Wikipedia, the free encyclopedia

In mathematics, **bilinear interpolation** is an extension of linear interpolation for interpolating functions of two variables on a regular grid. The key idea is to perform linear interpolation first in one direction, and then again in the other direction.

Suppose that we want to find the value of the unknown function f at the point $P = (x, y)$. It is assumed that we know the value of f at the four points $Q_{11} = (x_1, y_1)$, $Q_{12} = (x_1, y_2)$, $Q_{21} = (x_2, y_1)$, and $Q_{22} = (x_2, y_2)$.

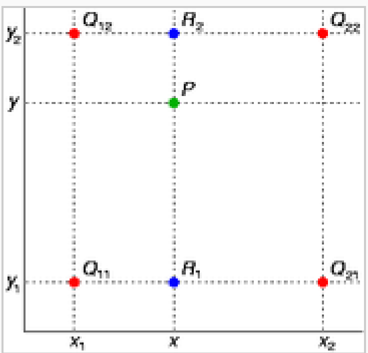

The four red dots show the data points and the green dot is the point at which we want to interpolate.

We first do linear interpolation in the x-direction. This yields

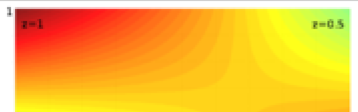

navigation
- Main page
- Contents
- Featured content
- Current events
- Random article

search
Go Search

interaction
- About Wikipedia
- Community portal
- Recent changes
- Contact Wikipedia
- Donate to Wikipedia
- Help

Bilinear interpolation

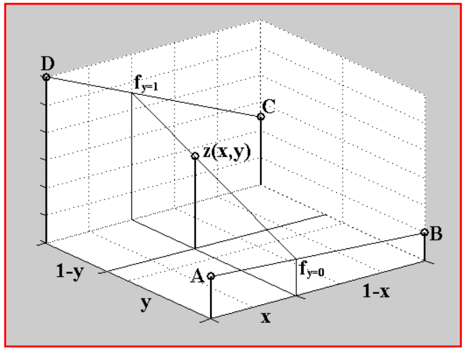

Bicubic polynomial interpolation

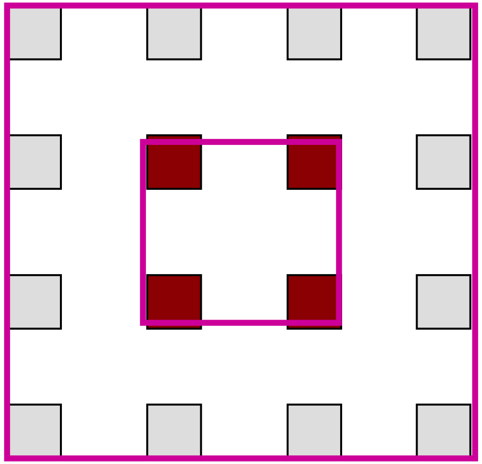

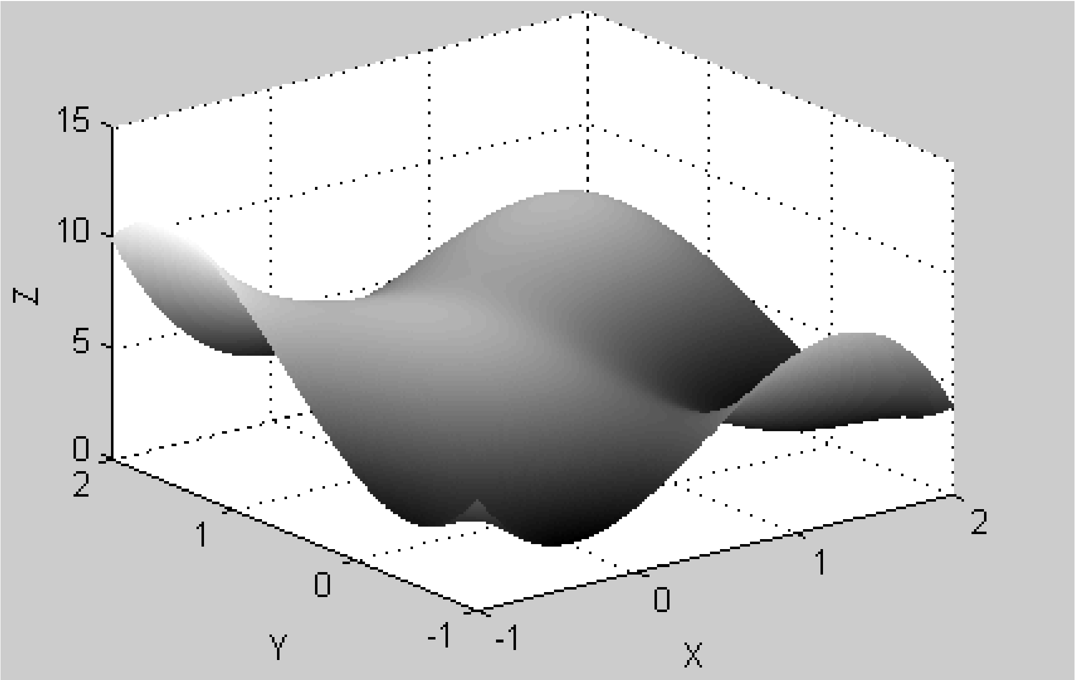

$$p(x,y)=\sum_{i=0}^{i=3}\sum_{j=0}^{j=3} a_{ij}x^i y^j$$

Bicubic interpolation - Wikipedia, the free encyclopedia - Windows Internet Explorer

http://en.wikipedia.org/wiki/Bicubic_interpolation

Google

Plik Edycja Widok Ulubione Narzędzia Pomoc

Bicubic interpolation - Wikipedia, the free encyclopedia

Strona Narzędzia

Log in / create account

article discussion edit this page history

WIKIPEDIA
The Free Encyclopedia

navigation
- Main page
- Contents
- Featured content
- Current events
- Random article

search

Go Search

interaction
- About Wikipedia
- Community portal
- Recent changes
- Contact Wikipedia
- Donate to Wikipedia
- Help

Bicubic interpolation

From Wikipedia, the free encyclopedia

In mathematics, **bicubic interpolation** is an extension of cubic interpolation for interpolating data points on a two dimensional regular grid. The interpolated surface is smoother than corresponding surfaces obtained by bilinear interpolation or nearest-neighbor interpolation. Bicubic interpolation can be accomplished using either Lagrange polynomials, cubic splines or cubic convolution algorithm.

In image processing, bicubic interpolation is often chosen over bilinear interpolation or nearest neighbor in image resampling, when speed is not an issue. Images resampled with bicubic interpolation are smoother and have fewer interpolation artifacts.

Contents [hide]

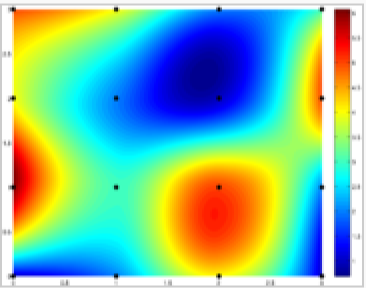

Bicubic interpolation on the square $[0,3] \times [0,3]$ consisting of 9 unit squares patched together. Bicubic interpolation as per Matlab's implementation. Colour

Bicubic spline interpolation [edit]

Suppose the function values f and the derivatives f_x, f_y and f_{xy} are known at the four corners $(0,0)$, $(1,0)$, $(0,1)$, and $(1,1)$ of the unit square. The interpolated surface can then be written

Bicubic polynomial interpolation

bicubic, bilinear, nearest neighbor (comparison)

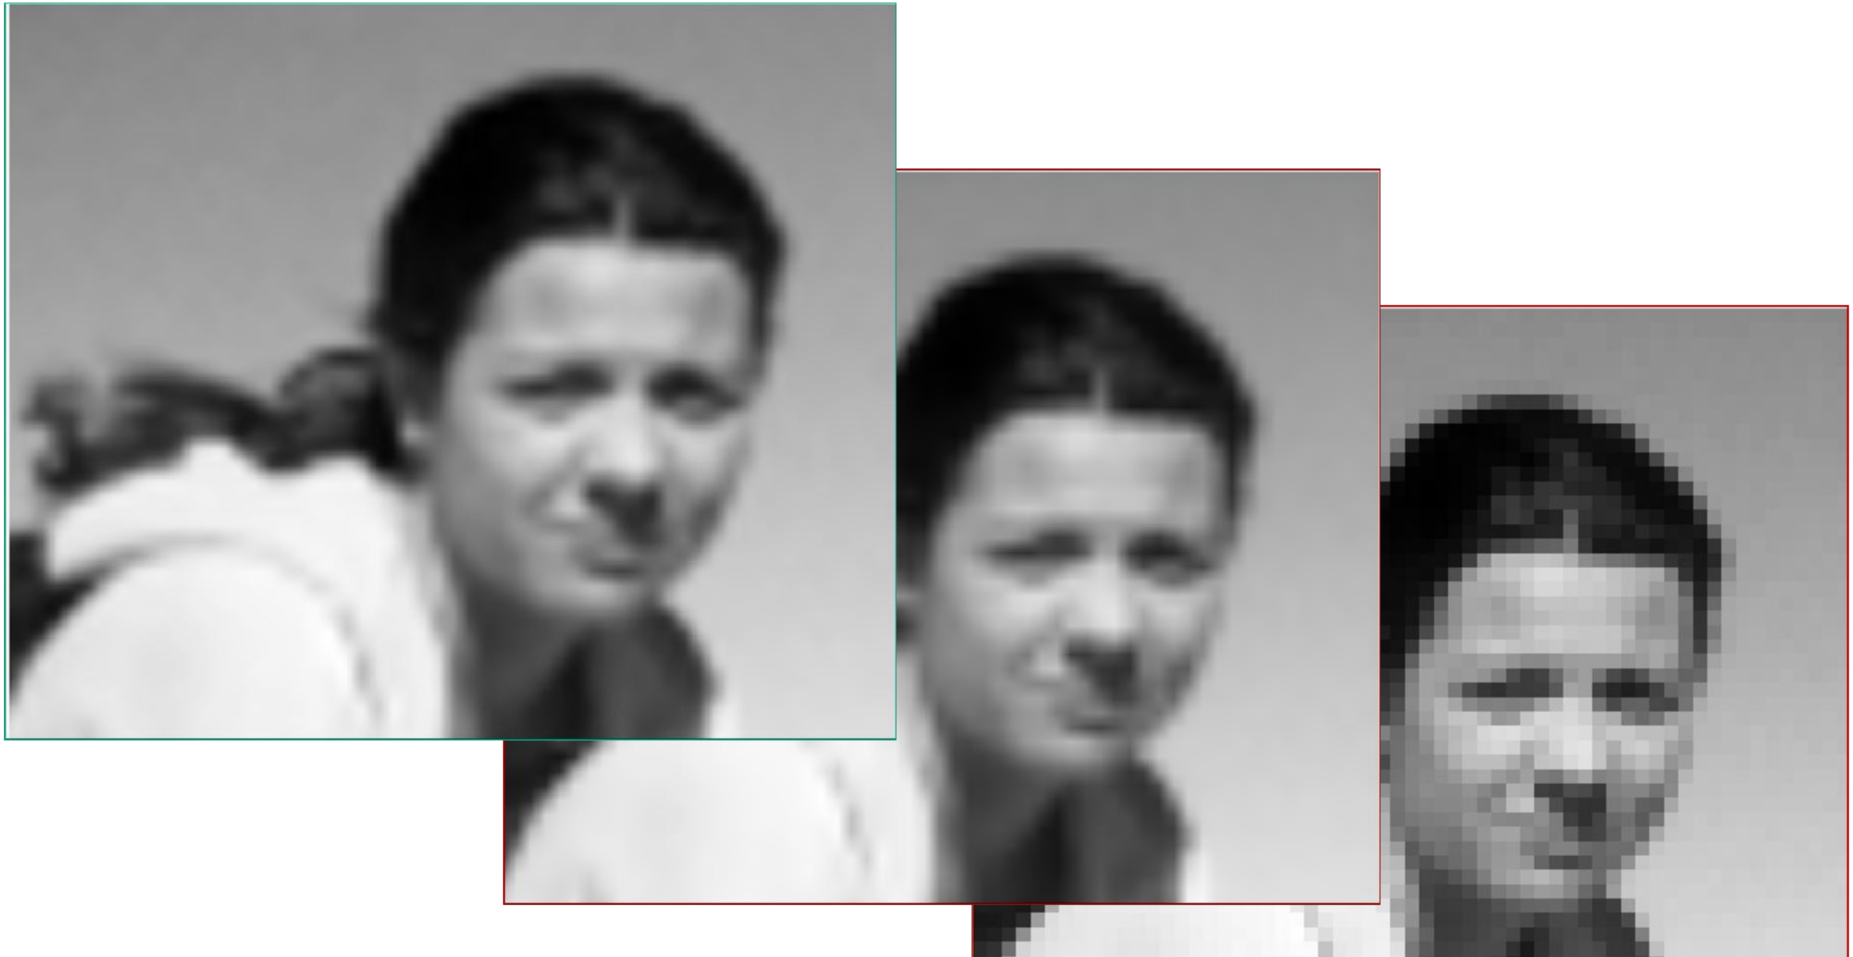

http://en.wikipedia.org/wiki/Bicubic_interpolation

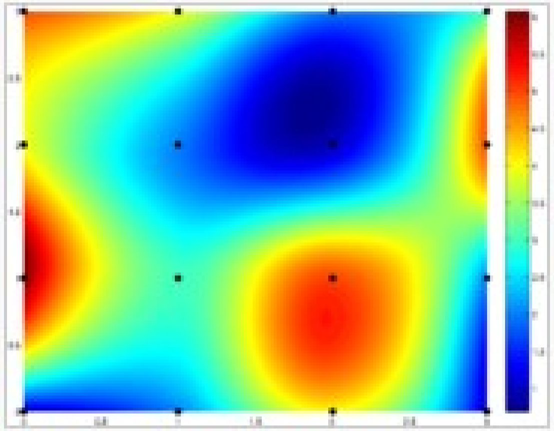

Bicubic interpolation on the square $[0,3] \times [0,3]$ consisting of 9 unit squares patched together. Bicubic interpolation as per Matlab's implementation. Colour indicates function value. The black dots are the locations of the prescribed data being interpolated. Note how the color samples are not radially symmetric. They are more square-based (It may be easier to

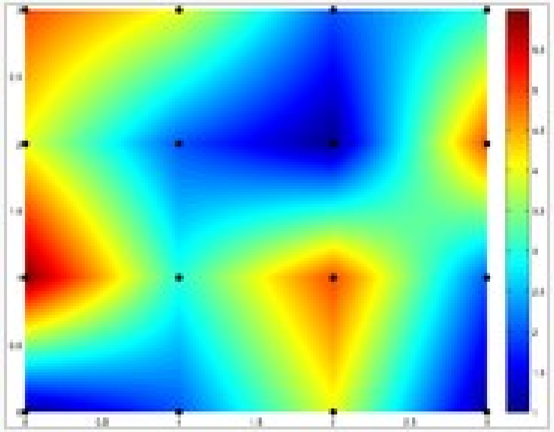

Bilinear interpolation on the same dataset as above. Derivatives of the surface are not continuous over the square boundaries.

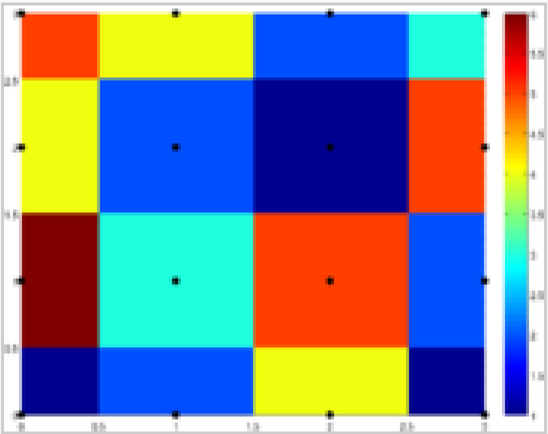

Nearest-neighbor interpolation on the same dataset as above. Note that the information content in all these three examples is equivalent.

Quiz

1. What is a size and shape of standard median filter?
2. What should be the size of a structuring element?
3. What is a difference between the sampling and the Nyquist frequency?
4. How many other pixels are used for interpolation of a pixel?
5.