

Computer Vision

MR2 Informatics Option GVR
James L. Crowley

Fall Semester

7 Nov 2019

Lesson 4

Local Image Description with Receptive Fields

Lesson Outline:

1. Describing Appearance with Local Features	2
1.1. Local Invariant descriptions of Appearance	2
1.2. Describing Local Appearance with Receptive Fields	3
1.3. The Mammalian Visual Cortex	4
1.4. Receptive Fields in the Visual Cortex.....	4
2. Using the Gaussian Function to Compute Image Derivatives	6
2.1. The Gaussian Function.....	6
2.2. Gaussian Derivative Operators	7
2.3. 2D Gaussian functions	7
2.4. Sampled Gaussian Derivative Filters in 1D.....	9
2.5. The 2D Sampled Gaussian Function.....	9
2.6. Using the Gaussian to compute image derivatives	10
3. Image Description Using Gaussian Derivatives.....	11
3.1. Multi-Resolution Image Gradient	11
3.2. Steerability of Gaussian Derivatives.....	12
3.3. Intrinsic Orientation	12
3.4. The Laplacian of the Gaussian and the DoG	13
4. Using the Gaussian function as a low-pass digital filter	14
4.1. Sampling (Optional advanced subject)	15
4.2. Setting the Window Size (Optional advanced subject).....	16

1. Describing Appearance with Local Features

1.1. Local Invariant descriptions of Appearance

For a computer, an image is just a meaningless table of numbers. To detect, track or recognize phenomena, we need to construct a description of contents in the image. Ideally this description should be local, invariant and multi-scale.

Information that can be perceived in an image is referred to as "phenomena". Phenomena may be found many scales in an image.

Scale and Resolution

Scale refers to the unit of measure to describe space. On a sheet of paper, we measure scale with centimeters. In a city, we measure position in kilometers. If we print a map of a city on a paper, we specify the "scale" of the map to tell how many cm represent a km.

Resolution is the ability to "resolve" detail. In a high-resolution map of the city, you can "resolve" small details such as sidewalks and entrances to buildings. In a low-resolution version, we can only resolve large structures such as roads and buildings.

Multi-scale vs Multi-resolution descriptions

A low-resolution map of a city will give global structure such as relative position of rivers, highways and neighborhoods. A low-resolution description provides context for high-resolution details such as streets and buildings. This shows the importance of description at multiple scales.

We will work with sampled images. The sample rate (distance between pixels) determines the scale of an image and sets a lower bound on the resolution. We will see in this lecture that we can oversample an image and obtain a low resolution image independent of scale.

Locality

Information (phenomena) appear in an image over a range of scales and positions. A position in an image and its immediate neighborhood are referred to as a "locality". The size of the neighborhood depends on both scale and resolution. Typically, the neighborhood is expressed in "pixels". The sample rate (or sample size, or pixel size) determines the scale of the neighborhood.

Today, we will look at the question: How many pixels should be used to represent a local image neighborhood. The number of samples determines the maximum information that can be represented.

Appearance

Appearance is what you see. For an image, appearance refers to the structures that we can describe in the image. We seek a description of appearance that will not change when the appearance changes in position, orientation, or scale. We refer to such a description as **invariant**.

Position, orientation and scale of appearance will change in predictable ways as a camera moves, or as objects moves. We seek a description that is position invariant (shift invariant), orientation invariant and scale invariant so that we obtain the same description independent of the relative position of the object.

1.2. Describing Local Appearance with Receptive Fields

Consider an image $P(i, j)$. Ideally we want to describe local appearance with a set of local functions, that describe the "appearance" around a point in an image. In the biological literature, these functions are called "receptive fields".

Ideally we should have a family of k such functions, $f_k(x, y)$ (x and y are integers). Each receptive field, $f_k(x, y)$ responds to some local pattern of appearance at an image position (i, j) by providing a feature value, $a_k(i, j)$

$$\bar{A}(i, j) = \begin{pmatrix} a_1 \\ a_2 \\ \dots \\ a_k \end{pmatrix}$$


The features values are computed by **convolution** of the receptive field functions with the image:

$$a_k(i, j) = f_k * P(i, j) = \sum_{x=-R}^R \sum_{y=-R}^R P(i-x, j-y) f_k(x, y)$$

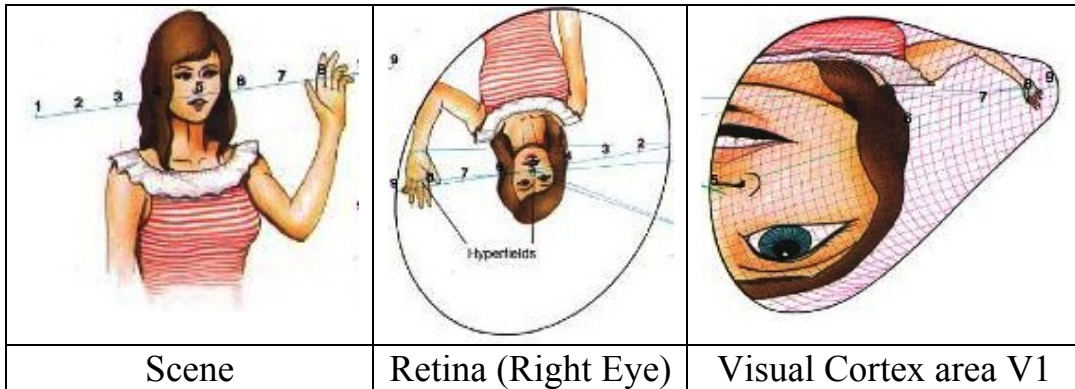
The value R is a radius that determines the locality for the feature. Note in this class, we will NEVER use $*$ for multiplication. We can find examples of such functions in the human visual cortex.

1.3. The Mammalian Visual Cortex

The Visual Cortex of mammals is composed of multiple layers of retinotopic maps. Each map is an image of the retina projected onto (convolved with) a receptive field at different scales and orientations.




The of Retino-topic maps occur over a range of scales and orientations





1.4. Receptive Fields in the Visual Cortex

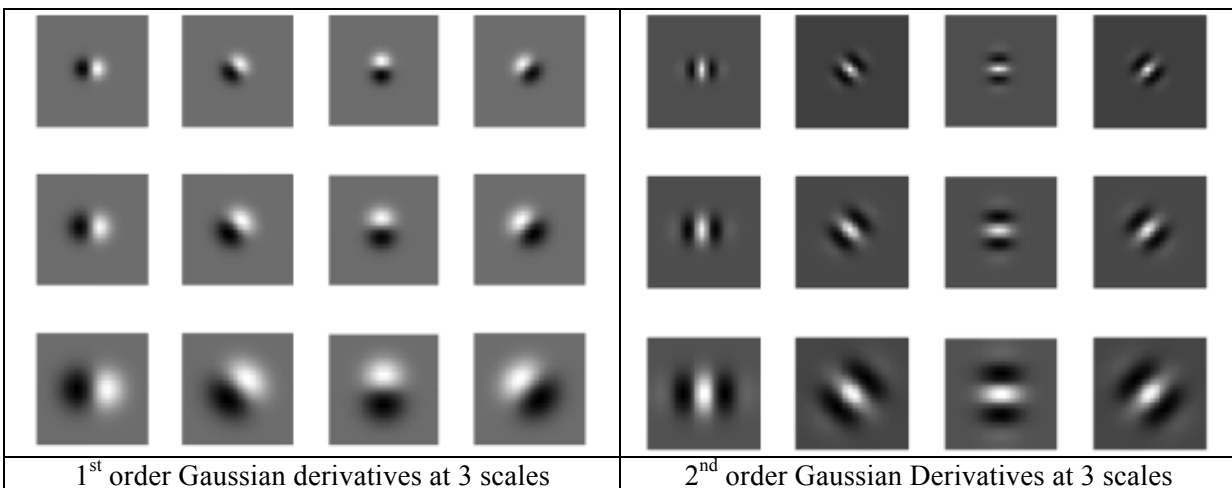
In 1968, Hubel and Wiesel probed the visual cortex of a cat with electrodes and found layers of cells that responded to local patterns of stimulation. They discovered that the visual cortex is composed of a series of layers. Each layer is a map of the retina filtered by a “receptive field” that respond to a certain pattern over a narrow range of sizes and orientations.



The patterns at the lowest level look like these:



Each layer is a specific pattern at a specific orientation and scale.



This figure shows first and second order Gaussian derivative features from 3 scales, computed using sums and differences of adjacent samples in a half-octave Gaussian pyramid.

The first level receptive fields were found to be local filters that pass a narrow range of spatial frequencies. The first layer receptive fields were found to be modeled by Gabor functions (Gaussian modulated by a Cosine plus an imaginary Sin).

The first layer receptive fields can also be modeled multi-scale derivatives of Gaussians functions. These two representations are very similar. However, Gabor functions can be very expensive to convolve with an image. Multi-scale Gaussian derivatives can convolved very efficiently because of a number mathematical properties that we will discuss below.

As they moved up the visual cortex, Hubel and Weisel found that these patters were combined to form more complex patterns, such as corners, bars, crosses, etc. These were named "complex" receptive field.

It is possible to learn receptive fields for local image description using back-propagation. This is useful for tuning a system to a specific problem domain. However, the general case of vision in the real world, this can require a very large set of training data and computing time.

Researchers in several different communities have found that as the variety of training data increases, the learned receptive fields converge to Gaussian Derivative Functions. This result has been demonstrated by researchers in both machine learning and computer vision, since the 1990's.

2. Using the Gaussian Function to Compute Image Derivatives

For a variety of reasons, derivatives of the Gaussian function emerge as general local image features. Chief among these are invariance to affine transformations.

In addition, using a well-defined mathematical function for receptive fields makes it possible to predict and explain system behavior, AND to avoid the very high cost of learning general functions. Gaussian functions also make possible a highly efficient computation for the lower levels of a deep network for image recognition.

In the next two lectures we will see how to use Gaussian derivatives to define general purpose receptive fields for image descriptions that are invariant to scale, orientation, and illumination, and robust to changes in viewing direction and illumination color.

2.1. The Gaussian Function

The Gaussian Function is $G(x, \sigma) = e^{-\frac{x^2}{2\sigma^2}}$

The standard deviation, σ , is often referred to as the scale of the Gaussian. Scale determines the size (spatial extent) of the locality of the description.

The Gaussian function is invariant to affine transformations.

$$T_a\{G(x, y, \sigma)\} = G(T_a\{x\}, T_a\{y\}, T_a\{\sigma\})$$

This invariance is a consequence of the fact that Gaussian functions are based on moments, and moments are affine invariant.

This is just one of the many interesting properties of the Gaussian function that make them very well suited as a basis function for local image description.

2.2. Gaussian Derivative Operators

The Gaussian function is: $G(x, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x^2}{2\sigma^2}}$

Fourier Transform: $F\{e^{-\frac{x^2}{2\sigma^2}}\} = G(\omega, \sigma) = \sqrt{2\pi} \sigma e^{-\frac{1}{2}\sigma^2\omega^2}$

Scale property: $G(x, \sqrt{2}\sigma) = G(x, \sigma) * G(x, \sigma)$

Derivatives:

$$\frac{\partial G(x, \sigma)}{\partial x} = -\frac{x}{\sigma^2} G(x, \sigma) = G_x(x, \sigma)$$

$$\frac{\partial^2 G(x, \sigma)}{\partial x^2} = \frac{x^2 - \sigma^2}{\sigma^4} G(x, \sigma) = G_{xx}(x, \sigma)$$

$$\frac{\partial^3 G(x, \sigma)}{\partial x^3} = \frac{x^3 - 3x\sigma^2}{\sigma^6} G(x, \sigma) = G_{xxx}(x, \sigma)$$

Note that the derivative of a Gaussian "contains" the original Gaussian.
Derivatives are modulated Gaussians!

2.3. 2D Gaussian functions

2D Gaussian Kernel: $G(x, y, \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}}$

Separability: $G(x, y, \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}} = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x^2}{2\sigma^2}} \cdot \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{y^2}{2\sigma^2}}$

But also $G(x, y, \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}} = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x^2}{2\sigma^2}} * \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{y^2}{2\sigma^2}}$

The convolution reduces to a product because the two components are orthogonal.
(Different free variables.)

This means that a convolution with a 2D Gaussian ($O(N^2)$ operations) can be computed a sequence of 2 convolutions with 1D Gaussians ($O(N)$ operations).

Scale property: $G(x, y, \sigma) * G(x, y, \sigma) = G(x, y, \sqrt{2}\sigma)$

In General $G(x, y, \sigma_1) * G(x, y, \sigma_2) = G(x, y, \sqrt{\sigma_1^2 + \sigma_2^2})$

The convolution of a two Gaussians yields a scaled Gaussian.

Derivatives:

$$\frac{\partial G(x,y,\sigma)}{\partial x} = -\frac{x}{\sigma^2} G(x,y,\sigma) = G_x(x,y,\sigma)$$

$$\frac{\partial G(x,y,\sigma)}{\partial y} = -\frac{y}{\sigma^2} G(x,y,\sigma) = G_y(x,y,\sigma)$$

$$\frac{\partial^2 G(x,y,\sigma)}{\partial x^2} = \frac{x^2 - \sigma^2}{\sigma^4} G(x,y,\sigma) = G_{xx}(x,y,\sigma)$$

$$\frac{\partial^2 G(x,y,\sigma)}{\partial x \partial y} = \frac{xy}{\sigma^4} G(x,y,\sigma) = G_{xy}(x,y,\sigma)$$

$$\frac{\partial^3 G(x,y,\sigma)}{\partial x^3} = \frac{x^3 - 3x\sigma^2}{\sigma^6} G(x,y,\sigma) = G_{xxx}(x,y,\sigma)$$

The Laplacian of the Gaussian: $\nabla^2 G(x,y,\sigma) = G_{xx}(x,y,\sigma) + G_{yy}(x,y,\sigma)$

The Diffusion Equation: $\nabla^2 G(x,y,\sigma) = \frac{\partial^2 G(x,y,\sigma)}{\partial x^2} + \frac{\partial^2 G(x,y,\sigma)}{\partial y^2} = \frac{\partial G(x,y,\sigma)}{\partial \sigma}$

As a consequence: $\nabla^2 G(x,y,\sigma) \approx (G(x,y,\sigma_1) - G(x,y,\sigma_2))$

This is called a Difference of Gaussian (DoG) and typically requires $\sigma_1 \geq 1.4 \sigma_2$
 It is common to use: $\nabla^2 G(x,y,\sigma) \approx G(x,y,\sqrt{2}\sigma) - G(x,y,\sigma)$


But note that from the scale property: $G(x,y,\sqrt{2}\sigma) \approx G(x,y,\sigma) * G(x,y,\sigma)$

so that $\nabla^2 G(x,y,\sigma) \approx G(x,y,\sigma) * G(x,y,\sigma) - G(x,y,\sigma)$

(note - this requires that $G(x,y,\sigma)$ to be normalized to sum to 1.

We can use these functions to create a basis set of receptive fields for appearance

$$G = (G_x, G_y, G_{xx}, G_{xy}, G_{yy}, G_{xxx}, G_{xxy}, G_{xyy}, G_{yyy})$$



The Gaussian receptive fields $G_x, G_y, G_{xx}, G_{xy}, G_{yy}, G_{xxx}, G_{xxy}, G_{xyy}, G_{yyy}$.

2.4. Sampled Gaussian Derivative Filters in 1D.

To work with images on a computer we need sampled Gaussians. We can do this by replacing the continuous variable x by a discrete integer n . Attention. THIS has consequences that we will see below!!!

The sampled Gaussian and its derivatives are:

$$G(n, \sigma) = e^{-\frac{n^2}{2\sigma^2}}$$

$$G_x(n, \sigma) = -\frac{n}{\sigma^2} G(n, \sigma) = -\frac{n}{\sigma^2} e^{-\frac{n^2}{2\sigma^2}}$$

$$G_{xx}(n, \sigma) = \frac{n^2 - \sigma^2}{\sigma^4} G(n, \sigma) = \frac{n^2 - \sigma^2}{\sigma^4} e^{-\frac{n^2}{2\sigma^2}}$$

$$G_{xxx}(n, \sigma) = -\frac{n^3 - n\sigma^2}{\sigma^6} G(n, \sigma) = -\frac{n^3 - n\sigma^2}{\sigma^6} e^{-\frac{n^2}{2\sigma^2}}$$

Note that there is only one parameter: σ . This determines the limit of the resolution for the position of a contrast point.

Note the scale parameter σ determines the "resolution" of the derivatives. You MUST specify σ . The smallest σ is not always the best. Many computer vision algorithms give unpredictable results because the researchers forget to specify the scale σ at which the algorithm was validated.

2.5. The 2D Sampled Gaussian Function

The 2D Gaussian Receptive Field is : $G(i, j, \sigma) = \frac{1}{B} W_N(i, j) \cdot e^{-\frac{(i^2+j^2)}{2\sigma^2}}$

where

$$w_N(i, j) = \begin{cases} 1 & \text{for } -R \leq i \leq R \text{ and } -R \leq j \leq R \\ 0 & \text{otherwise} \end{cases}$$

Finite window, $w_N(i, j)$ has $N^2 = (2R+1)^2$ coefficients

Typically: for R should be $\geq 3\sigma$. Recommend $R=4\sigma$

The normalization factor $B = \sum_{x=-R}^R \sum_{y=-R}^R e^{-\frac{(i^2+j^2)}{2\sigma^2}} \approx 2\pi\sigma$

2.6. Using the Gaussian to compute image derivatives

For an image $p(i,j)$, the derivatives can be approximated by convolution with Derivatives of Gaussians.

$$\frac{\partial p(i,j)}{\partial x} * G(x,y) = \frac{\partial}{\partial x} * p(i,j) * G(x,y) = \frac{\partial}{\partial x} * G(x,y) * p(i,j) = \frac{\partial G(x,y)}{\partial x} * p(i,j)$$

Thus we can approximate an image derivative as $p_x(i,j) \approx G_x * p(i,j)$

By substituting (i,j) for (x,y) . (Note that the derivative order is x).

$$p_x(i,j) \approx p * G_x(i,j) = p(i,j) * G_x(i,j) = G_x * p(i,j)$$

IMPORTANT: To compute G_x , it is **NECESSARY** to specify σ . This determines the "resolution" of the derivative.

Small σ is not necessarily best.

$$p_x(i,j,\sigma) \approx G_x(x,y,\sigma) * p(i,j) = G_x(\sigma) * p(i,j)$$

Similarly: $p_y(i,j,\sigma) \approx G_y(\sigma) * p(i,j)$

$$p_{xx}(i,j,\sigma) \approx G_{xx}(\sigma) * p(i,j)$$

$$p_{xy}(i,j,\sigma) \approx G_{xy}(\sigma) * p(i,j)$$

$$p_{yy}(i,j,\sigma) \approx G_{yy}(\sigma) * p(i,j)$$

The Gradient of the image $\vec{\nabla}p(i,j)$ is calculated by $\vec{\nabla}G(\sigma) * p(i,j)$

where $\vec{\nabla}G(\sigma) = \begin{pmatrix} G_x(\sigma) \\ G_y(\sigma) \end{pmatrix}$ This gives:

$$\text{Gradient: } \vec{\nabla}p(i,j,\sigma) = \begin{pmatrix} p_x(i,j,\sigma) \\ p_y(i,j,\sigma) \end{pmatrix} \approx \vec{\nabla}G(\sigma) * p(i,j) = \begin{pmatrix} G_x(\sigma) \\ G_y(\sigma) \end{pmatrix} * p(i,j)$$

Laplacien:

$$\nabla^2 p(i,j,\sigma) = \nabla^2 G(\sigma) * p(i,j) = p_{xx}(i,j,\sigma) + p_{yy}(i,j,\sigma) \approx G_{xx}(\sigma) * p(i,j) + G_{yy}(\sigma) * p(i,j)$$

To use Gaussian functions to describe images we need to sample the Gaussian and limit its extent. That is, we must define Gaussian Filters.

3. Image Description Using Gaussian Derivatives

3.1. Multi-Resolution Image Gradient

The Gradient of the image $\vec{\nabla}p(i, j)$ is calculated by $\vec{\nabla}G(\sigma) * p(i, j)$

where
$$\vec{\nabla}G(\sigma) = \begin{pmatrix} G_x(\sigma) \\ G_y(\sigma) \end{pmatrix}$$

The gradient of the image is:
$$\vec{\nabla}p(i, j, \sigma) = \begin{pmatrix} p_x(i, j, \sigma) \\ p_y(i, j, \sigma) \end{pmatrix} \approx \vec{\nabla}G(\sigma) * p(i, j) = \begin{pmatrix} G_x(\sigma) * p(i, j) \\ G_y(\sigma) * p(i, j) \end{pmatrix}$$

The second derivative of the image (the Laplacien) is:

$$\nabla^2 p(i, j, \sigma) = \nabla^2 G(\sigma) * p(i, j) = p_{xx}(i, j, \sigma) + p_{yy}(i, j, \sigma) \approx G_{xx}(\sigma) * p(i, j) + G_{yy}(\sigma) * p(i, j)$$

3.2. Steerability of Gaussian Derivatives.

It is possible to synthesize an oriented derivative at any point as a weighted sum of derivatives in perpendicular directions. The weights are given by sine and cosine functions. The weights are given by sine and cosine functions.

$$G_x^\theta(x, y, \sigma) = \cos(\theta) \cdot G_x(x, y, \sigma) + \sin(\theta) \cdot G_y(x, y, \sigma)$$

Higher order derivatives can also be steered.

Thus:

$$\text{1st order} \quad p_x^\theta(i, j, \sigma) = \text{Cos}(\theta)p_x(i, j, \sigma) + \text{Sin}(\theta)p_y(i, j, \sigma)$$

$$\text{2nd order} \quad p_{xx}^\theta(i, j, \sigma) = \text{Cos}(\theta)^2 p_{xx}(i, j, \sigma) + 2\text{Cos}(\theta)\text{Sin}(\theta)p_{xy}(i, j, \sigma) + \text{Sin}(\theta)^2 p_{yy}(i, j, \sigma)$$

3rd order

$$p_{xxx}^\theta(i, j, \sigma) = \text{Cos}(\theta)^3 p_{xxx}(i, j, \sigma) + 3 \cdot \text{Cos}(\theta)^2 \text{Sin}(\theta)p_{xxy}(i, j, \sigma) + 3 \cdot \text{Cos}(\theta)\text{Sin}(\theta)^2 p_{xyy}(i, j, \sigma) + \text{Sin}(\theta)^3 p_{yyy}(i, j, \sigma)$$

By steering the derivatives to the local orientation, we obtain an "invariant" measure of local contrast. We can also "steer" in scale to obtain invariance to size.

Note, we can NOT steer the mixed derivatives, i.e. $p_{xy}(i, j, \sigma)$

3.3. Intrinsic Orientation

For each pixel, one can calculate the orientation of maximal gradient. This orientation is equivariant with rotation. One can use this as an "intrinsic" orientation to normalize the receptive fields at any point in the image.

$$\text{Local orientation: } \theta_i(x, y, \sigma) = \text{Tan}^{-1}\left(\frac{G_y \cdot P(x, y, \sigma)}{G_x \cdot P(x, y, \sigma)}\right)$$

Note that local orientation depends on σ !

3.4. The Laplacian of the Gaussian and the DoG

The Laplacian of Gaussian is a scalar value:

$$\nabla^2 G(x,y,\sigma) = G_{xx}(x,y,\sigma) + G_{yy}(x,y,\sigma) = \frac{\partial G(x,y,\sigma)}{\partial \sigma}$$

Because it is the derivative with respect to σ , it can be approximated by a difference of Gaussians (DoG):

$$\nabla^2 G(x,y,\sigma) \approx G(x,y,\sigma_1) - G(x,y,\sigma_2)$$

This is called a Difference of Gaussian and typically requires $\sigma_1 \geq 1.4 \sigma_2$


It is common to use: $\nabla^2 G(x,y,\sigma) \approx G(x,y,\sqrt{2}\sigma) - G(x,y,\sigma)$

Because of the scale property: $G(x,y,\sqrt{2}\sigma) \approx G(x,y,\sigma) * G(x,y,\sigma)$

We can easily compute a DoG as

$$\nabla^2 G(x,y,\sigma) \approx G(x,y,\sigma) * G(x,y,\sigma) - G(x,y,\sigma)$$

In 1D:



4. Using the Gaussian function as a low-pass digital filter

(Optional advanced subject)

Computers represent image as 2D sampled digitized signals. Because they are sampled, processing requires convolution with a sampled filter.

To verify the properties of a Digital (sampled) Gaussian we will need its Fourier Transform:

Fourier Transform:
$$F\left\{e^{-\frac{x^2+y^2}{2\sigma^2}}\right\} = \frac{\pi}{2\sigma^2} e^{-\frac{1}{2}\sigma^2(u^2+v^2)}$$

To obtain a digital Gaussian filter we must perform two operations:

- 1) Sample the spatial axis x, y at a rate of Δx , and Δy
- 2) Limit the spatial extent with a window $W_N(x,y)$

$$G(x,y;\sigma) \rightarrow G(i,j;\sigma) \equiv W_N(i,j)G(i\Delta x, j\Delta y;\sigma)$$

Thus there are 2 parameters to Control:

- 1) Sample Distance Δx
- 2) Window size, $N = 2R+1$

These are both determined by “scale” parameter of the Gaussian: σ

Sample Distance: Easy answer – Let $\Delta x = 1$ and control σ .

This is valid, provided that $\sigma \geq \Delta x$ or that $\sigma/\Delta x \leq 1$

Window Size: $R \geq 3\sigma$ Thus $N \geq 6\sigma+1$

Note that $R = 3\sigma$ is a lower limit that can leave some windowing noise in the function.

4.1. Sampling (Optional advanced subject)


Let us consider the case of a 1-D Gaussian.

$$G(x, \sigma) = e^{-\frac{x^2}{2\sigma^2}}$$

To sample we replace x with $n\Delta x$.

$$G(n\Delta x, \sigma) = e^{-\frac{(n\Delta x)^2}{2\sigma^2}}$$

This is modeled as multiplication by an infinite pulse chain.



where:


$$\delta_{\Delta x}(x) = \sum_{n=-\infty}^{\infty} \delta_{\Delta x}(n\Delta x)$$

So that

$$G(n) = G(x) \cdot \Delta x \delta_{\Delta x}(x) = \sum_{n=-\infty}^{\infty} G(x) \cdot \delta_{\Delta x}(x - n\Delta x)$$

Multiplication in Space is a Convolution in Frequency. The Fourier transform of the sampling function is:

$$F(\delta_{\Delta x}(x)) = \Delta x \sum_{n=-\infty}^{\infty} \delta(n f_{\Delta x})$$



The ideal sample function is a

Where $f_{\Delta x}$ is the "Nyquist" frequency $f_{\Delta x} = \frac{1}{2\Delta x}$

In the Frequency domain, sampling converts the Fourier Transform of the Gaussian into an infinite sequence of Gaussians.

Transform of the Gaussian is

$$F\left\{e^{-\frac{x^2}{2\sigma^2}}\right\} = G(\omega, \sigma) = \sqrt{2\pi} \sigma e^{-\frac{1}{2}\sigma^2\omega^2}$$

Sampling creates multiple copies intervals of $G(\omega)$ at intervals of $f_{\Delta x} = \frac{1}{2}\Delta x$

The tail of the Gaussian beyond $f_{\Delta x} = \frac{1}{2}\Delta x$ will be converted to noise. We need to insure that the integral from f_n to infinite is small.

Rule of thumb: assure that $\sigma \geq \Delta x$

We can define the sample size to be $\Delta x = l$. This gives a sampled function

$$G(n, \sigma) = e^{-\frac{n^2}{2\sigma^2}}$$

4.2. Setting the Window Size (Optional advanced subject)

To represent this in a computer we must also specify the spatial extent (number of samples), N of the filter. We set $N = 2R + 1$ where R is the "radius" of the function.

This gives us 2 parameters to control:

- 1) The scale of the Gaussian $\sigma/\Delta x$
- 2) the size of the support $N = 2R+1$

Truncating a function to a finite support is equivalent to multiply by a window $W_N(n)$


When we limit $G(x, \sigma)$ to a finite support, we multiply by a window

$$G(n, \sigma) = G(n, \sigma) \cdot w_N(n) \text{ where } w_N(n) = \begin{cases} 1 & \text{for } -R \leq n \leq R \\ 0 & \text{otherwise} \end{cases}$$

(note $N = 2R+1$). Multiplying by a finite window is equivalent to convolving with the Fourier transform of the finite window:

$$F\{G(n, \sigma) \cdot w_N(n)\} = G(\omega, \sigma) * W_N(\omega)$$

where $W_N(\omega) = \frac{\sin(\omega N/2)}{\sin(\omega/2)}$ and $G(\omega, \sigma) = \sqrt{2\pi} \sigma e^{-\frac{1}{2}\sigma^2 \omega^2}$



For $N < 7$, the ripples in $W_N(\omega)$ dominate the spectrum and corrupt the resulting Gaussian.

At $N=7$ the effect is tolerable but significant.

At $N \geq 9$ the effect becomes minimal

In addition for $\sigma/\Delta x < 1$, the phenomenon of aliasing folds a significant amount of energy at the Nyquist frequency, corrupting the quality (and the invariance) of the Gaussian function.

Finally, it is necessary to assure that the "gain" of the Gaussian filter is 1. This can be assured by normalizing so that the sum of the coefficients is 1. If the Gaussian were infinite in extent, then

$$\sum_{x=-\infty}^{\infty} e^{-\frac{x^2}{2\sigma^2}} = \sqrt{2\pi}\sigma$$

However, because we truncate the Gaussian to an size $n \pm R$, we must calculate the sum of the coefficients, A:

$$A = \sum_{n=-R}^R e^{-\frac{n^2}{2\sigma^2}}$$

The Gaussian filter is thus normalized by dividing by A to give a unit gain Receptive Field.

$$G(n, \sigma) = \frac{1}{A} e^{-\frac{n^2}{2\sigma^2}}$$