Computer Vision
MR2 Informatics Option GVR
James L. Crowley

Fall Semester   13 November 2014
Lesson 5

Describing Local Appearance with Gaussian Derivatives

Lesson Outline:

1. Beyond Edges - Describing Local Appearance

Appearance is what you see. Consider an image p(i, j).

We seek a set of K local basis functions, $f_k(x,y)$ to describe "appearance" around a point in an image. These are referred to as "local image description functions" or simply "Local features". (x and y are integers).

Each function, $f_k(x,y)$ gives an image "feature", a_k, describing appearance in the neightborhood of the image position p(i, j).

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

Projection of the image neighborhood *p(i,j)* onto this set of functions gives a

"feature" vector for appearance, $\bar{A}(i,j) = \begin{pmatrix} a_1 \\ a_2 \\ ... \\ a_K \end{pmatrix}$ around position *p(i,j)*.

Ideally the set of local image description functions should be orthogonal

$$\sum_{x=-R}^{R} \sum_{y=-R}^{R} p(i-x, j-y) f_k(x,y) = 0 \quad \text{if } n \neq m \quad \text{for all } (i, j)$$

to minimize the number of features. However, it can be shown that a function cannot be both orthogonal and shift invariant. Because shift invariance is important for robust view invariant recognition, we compromise on orthoganality.

For a variety of reasons, derivatives of the Gaussian function have been found to be very useful as local image features. Chief among these are invariance to affine transformations.

2. Image Description Using Gaussian Derivatives

2.1. The Gaussian Function

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

The Gaussian function is invariant to affine transformations.

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

For example, a change in scale is an affine transform:

$$T_s\{G(x,\ y,\ \sigma)\}\ =\ G(T_s\{x\},\ T_s\{y\}, T_s\{\sigma\}) = G(sx,\ sy,\ s\sigma)$$

This is just one of the many interesting properties of the Gaussian function when used as the basis for an image descriptor.

2.2. Gaussian Derivatives Operators

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

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

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

Derivatives:

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

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

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

2.3. 2D Gaussian functions

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

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

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

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

Derivatives:
$$\frac{\partial G(x,y,\sigma)}{\partial x} = -\frac{x}{\sigma^2} G(x,y,\sigma) = G_x(x,y,\sigma)$$
$$\frac{\partial G(x,y,\sigma)}{\partial y} = -\frac{y}{\sigma^2} G(x,y,\sigma) = G_y(x,y,\sigma)$$
$$\frac{\partial^2 G(x,y,\sigma)}{\partial x^2} = \frac{x^2-\sigma^2}{\sigma^4} G(x,y,\sigma) = G_{xx}(x,y,\sigma)$$
$$\frac{\partial^2 G(x,y,\sigma)}{\partial x \partial y} = \frac{xy}{\sigma^4} G(x,y,\sigma) = G_{xy}(x,y,\sigma)$$
$$\frac{\partial^3 G(x,y,\sigma)}{\partial x^3} = \frac{x^3-x\sigma^2}{\sigma^6} G(x,y,\sigma) = G_{xxx}(x,y,\sigma)$$

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

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 required $G(x,y,\sigma)$ to be normalized to sum to 1.

4

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

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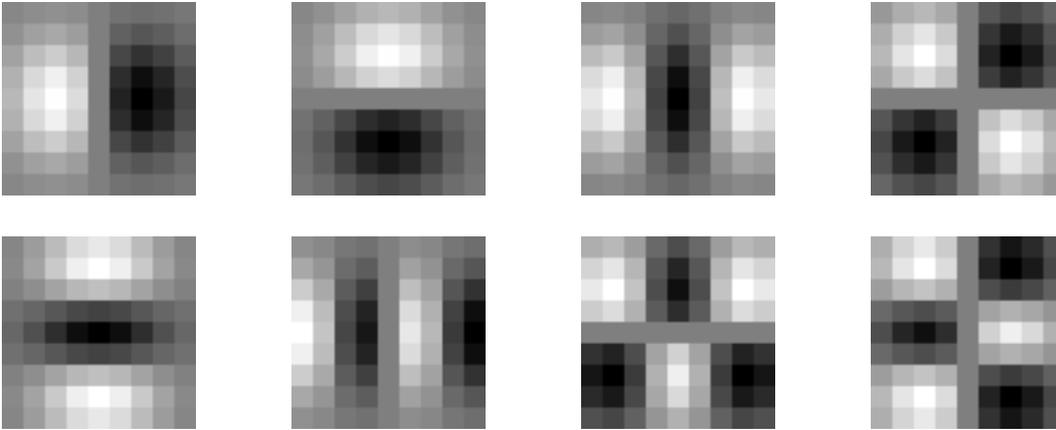
The Gaussian receptive fields $G_x, G_y, G_{xx}, G_{xy}, G_{yy}, G_{xxx}, G_{xxy}, G_{xyy}, G_{yyy}$.

2.4. A vector for describing local appearance.

Derivatives of the Gaussian function have been found to be very useful as local image features.   Chief among these are invariance to affine transformations.

To describe the local appearance in an image p(i, j), we "project" a local neighborhood (window) onto a set of feature functions, $G_k(x,y,\sigma)$.
These are referred to as "local image description functions" or "Local features".

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

For example:

$$\vec{G}^\sigma = G_k^\sigma = (G_x, G_y, G_{xx}, G_{yy}, G_{xy})$$

Note that you must specify σ. $\sigma=1$ is not necessarily best.

Each function, $G_k(x,y,\sigma)$ gives an image "feature", a_k, describing appearance in the neighborhood of the image position p(i, j). (let x and y be integers).

$$a_k(i,j) = \sum_{x=-R}^{R} \sum_{y=-R}^{R} p(i-x, j-y) G_k^\sigma(x,y)$$

Projection of the image neighborhood *p(i,j)* onto this set of functions gives a

"feature" vector for appearance, $\bar{A}(i,j) = \begin{pmatrix} a_1 \\ a_2 \\ ... \\ a_K \end{pmatrix}$ at each pixel *p(i,j)*.

We can use this feature vector much as we used color features - to detect objects based on their appearance.

2.5. 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 $P_x(i,j) \approx G_x * P(i,j)$
However to compute G_x, it is NECESSARY to specify σ.
Small σ is not necessarily best.

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

or more simply $p_x(i,j,\sigma) \approx G_x(\sigma) * p(i,j)$

Simarly:
$$p_y(i,j,\sigma) \approx G_y(\sigma) * p(i,j)$$
$$p_{xx}(i,j,\sigma) \approx G_{xx}(\sigma) * p(i,j)$$
$$p_{xy}(i,j,\sigma) \approx G_{xy}(\sigma) * p(i,j)$$
$$p_{yy}(i,j,\sigma) \approx G_{yy}(\sigma) * p(i,j)$$

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

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

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

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 an limit its extent. That is, we must define Gaussian Filters.

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

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

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

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.

3.1. 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.

$$G(x) \longrightarrow \otimes \longrightarrow G(n\Delta x)$$

$$\Delta x \, \delta_{\Delta x}(x)$$
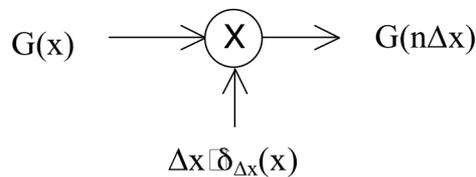
where:

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

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:

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

The ideal sample function is a

$$X(f)$$

$$-2f_e \quad -f_e \quad 0 \quad f_e \quad 2f_e$$
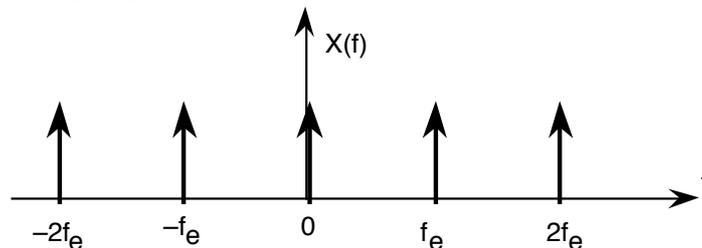
Where $f_{\Delta x}$ is the "Nyquist" frequency $f_{\Delta x} = \dfrac{1}{2\Delta x}$

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

Transform of the Gaussian is

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

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 $\Delta x = 1$. This gives a sampled function

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

3.2. 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 $\quad W_N(\omega) = \dfrac{\sin(\omega N/2)}{\sin(\omega/2)}\quad$ and $\quad G(\omega,\sigma) = \sqrt{2\pi}\ \sigma\, e^{-\frac{1}{2}\sigma^2\omega^2}$
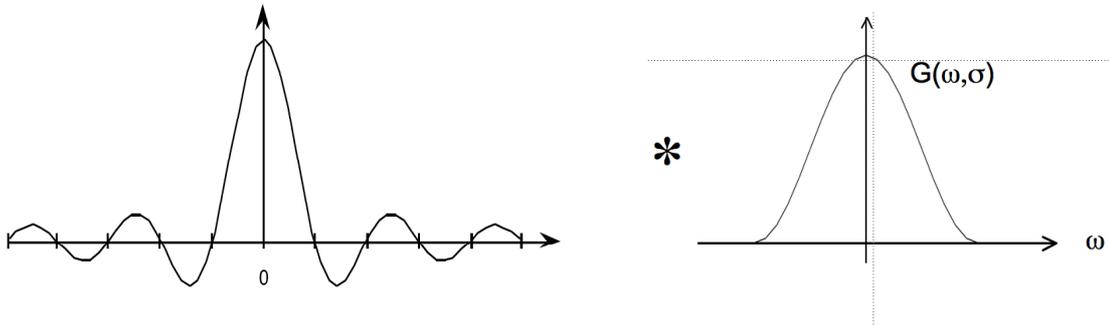
For N < 7, the ripples in WN(w) dominate the spectrum and corrupt the resulting Gaussian.

At N=7 the effect is tolerable but significant.

At N≥ 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 ±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.

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

3.3. 1-D Sampled Gaussian Derivative Filters

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.

3.4. 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

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

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

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

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

4. Using the Gaussian to compute image derivatives

For an image *p(i, j)*, the derivative can be approximated by convolution with the derivatives of a Gaussian.

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

Where $G(\sigma) = G(i,j,\sigma)$.

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

However to compute G_x, it is NECESSARY to specify σ.

Small σ is not necessarily best. Information exists at ALL values of σ.

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

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:

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

Laplacien:

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

4.1. 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:

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

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

3rd order

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

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 $p_{xy}(i, j, \sigma)$

4.2. 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.

Local orientation: $\quad \theta_i(x,y,\sigma) = Tan^{-1}(\dfrac{G_y \cdot P(x,y,\sigma)}{G_x \cdot P(x,y,\sigma)})$

Note that local orientation depends on σ!

5. Image Scale Space
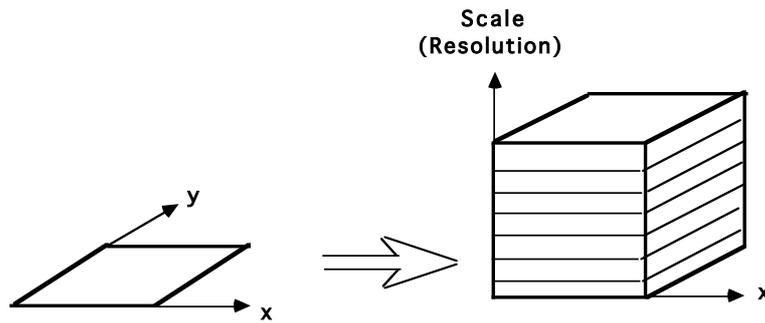
Scale space represents a 2D image in a 3D space, where the 3rd dimension is "resolution" or scale. This can be used to make image descriptions invariant to changes in size.

5.1. Continuous Scale Case.

Let p(x, y) be the image.
Let $G(x, y, \sigma)$ by a Gaussian function of scale $\sigma = 2^{s/2}$

The image Scale Space is a 3D continuous space *p(x, y, s)*

$$p(x, y, s) \; = \; p(x,y)* \, G(x, y, 2^s)$$

Note that the scale (s) axis is logarithmic. $s = Log_2(\sigma)$

5.2. Image Pyramids - Discrete Scale Space

Let *p(i,j)* be image of size N x N pixels.
We can create a discrete scale space by convolving *p(i,j)* with a set of Gaussians filters of increasing σ.

$$p(i, j, k) \; = \; p(i, j)* \, G(i, y, \sigma_k)$$

Scale invariance demands that the scale values, σ_k, should be from an exponential sequence. For example, we can sample the scale axis with a step size of $\Delta\sigma = 2$
In this case:

$$\sigma_k = 2^k$$

Because the Gaussian, $G(x, y, \sigma_k)$, is a low pass filter, as σ_k grows it becomes possible to resample the image with a larger step size without loss of information. Such resampling has the benefit of assuring an <u>invariance</u> of the impulse response of each image. The sample size Δx_k, Δy_k can grow exactly as σ_k.

$$p(i, j, k) = p(x \, \Delta x_k, y \, \Delta x_k, k)$$

What sample size is possible? It is possible to show that the sample step must be smaller than σ. For example, let $\Delta x = 2\sigma$

If we choose k such that $\Delta x_k = 2^k$, then $\sigma_k = 2^k$ for $k = 0$ to $Log_2(M)$

This defines a resampled scale space (a pyramid), with a scale step of one "octave".

$$p(i, j, k) = p(x \, \Delta x_k, y \, \Delta x_k, k) \quad \text{such that } \Delta x_k = 2^k \text{ and } \sigma_k = 2^k$$

This can be drawn as a set of images:
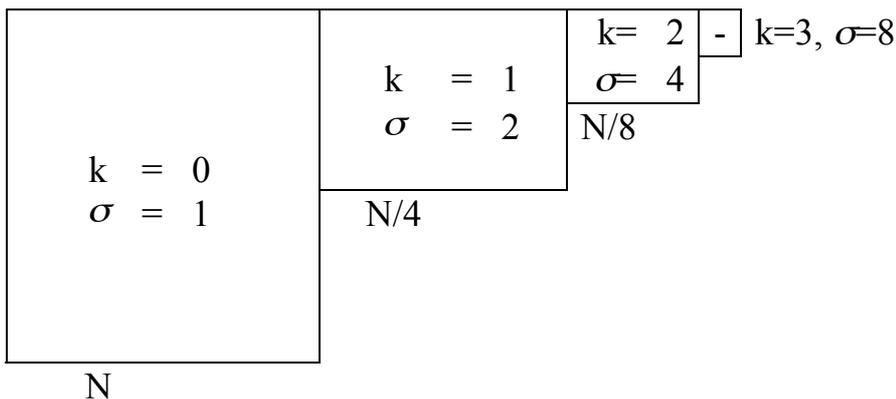
Let the image have M=NxN pixels. The total number of pixels is

$$P = M+M/4+M/8+M/16+M/64+...= 1.333... \text{ Pixels.}$$

When $\sigma/\Delta x$ is held constant, every "level" k of p(i, j, k) has the same impulse response.
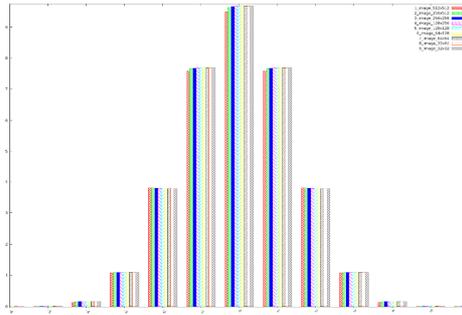
In fact, a step size of $\sigma_k = 2^k$ is too large. We need at least $\Delta\sigma_k = 2^{1/2}$
A more reasonable pyramid is given by

$$\sigma_k = 2^{(k+1)/2} \text{ and } \Delta x = 2^{k/2}$$

Example:

This involves some tricky resampling, but gives a pyarmid of size

$$P = M + M/2 + M/4 + M/8 + \ldots = 2M \text{ samples,}$$

5.3.   Laplacian Profile

At every image point, the Laplacian profile is the Laplacian of the image computed over a continuous (exponential) range of scales.

$$L(x, y, s) = P(x, y) * \nabla^2 G(x, y, 2^s)$$

18

The Laplacian profile is invariant to rotation and translation and equivariant to changes in scale. Since scale is proportional to distance, the profile is invariant to viewing distance.
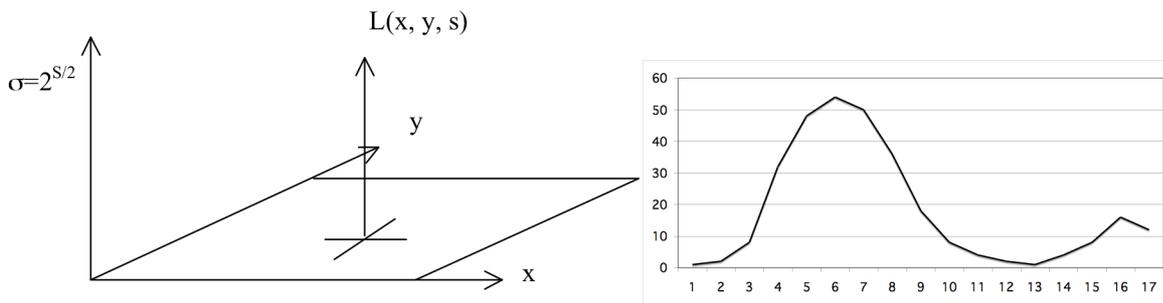

A change in viewing distance at x, y shifts the function L(x,y,s) in s.
The function remains the same. Thus the maximum is a local invariant.

The "intrinsic" scale at a point (x, y) is $\sigma_i = 2^{s_i}$

such that $s_i = \arg-\max_{s}\{L(x,y,s)\}$

Examples:
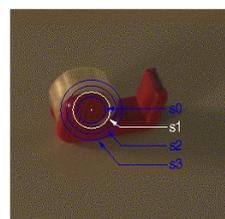
zero crossing of Laplacian at si
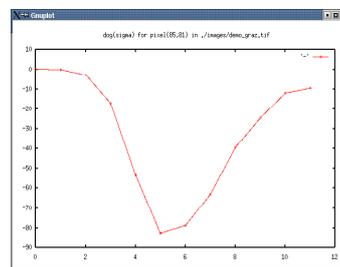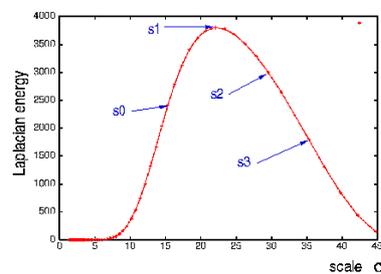
The scale of the maximal Laplacian is an invariant at ALL image points.