Discrete Scale Spaces via Heat Equation

ANDERSON CUNHA, RALPH TEIXEIRA AND LUIZ VELHO

IMPA–Instituto de Matemática Pura e Aplicada,
{mayrink|ralph|lvelho}@visgraf.impa.br

Abstract. Scale spaces allow us to organize, compare and analyse differently sized structures of an object. The linear scale space of a monochromatic image is the solution of the heat equation using that image as an initial condition. Alternatively, this linear scale space can also be obtained applying Gaussian filters of increasing variances to the original image. In this work, we compare (by looking at theoretical properties, running time and output differences) five ways of discretizing this Gaussian scale-space: sampling Gaussian distributions; recursively calculating Gaussian approximations; using Splines; approximating by first-order generators; and finally, by a new method we call "Crossed Convolutions". In particular, we explicitly present a correct way of initializing the recursive method to approximate Gaussian convolutions.

1 Introduction

The concept of *scale* is intrinsic to the observation of any physical object. In order to obtain information or features from an object, we must use an appropriate scale where such information is easily observable. For example, in cartography, if we want to find a house in a neighborhood, the scale 1:10.000 is appropriate. However, if we want the location of this same spot in the planet, we have to use a 1:1.000.000 scale.

An appropriate choice of scale is also imperative to locate phenomena in time. Certain physical phenomena are observable in time scales of micro-seconds (think of the information flow inside a computer), while others occur in time scale of billions of years (like the life cycle of stars).

Note that the concept of scale is closely linked to the idea of *resolution* (or detail). In general, the resolution is the inverse of scale: whenever we examine events of small (fine) scale, we use high resolutions.

If we want a generic system that does not know *a priori* the desired scale for extracting information from an object (or sequence of images), then it might be necessary to create a representation of such an object in several scales. In image processing, several such *multi-scale representations* have been created to face this difficulty, like Quad-Trees, Pyramids, Wavelets and Scale-Spaces.

Some desirable properties for a multi-scale representation are:

- Isotropy or rotation invariance: there are no preferred directions, that is, the representation of a structure is independent of its orientation.
- Homogeneity or translation invariance: there are no preferred locations, that is, the representation of a structure does not depend on its localization.
- Causality: no structure is created when the scale increases. More specifically, any feature present in a certain scale comes from details that are present in smaller scales; the representation of an object in a coarser scale has no more details than the same object represented on a smaller (finer) scale.

Note that these properties (specially the invariance by translations) imply that the size of the support of an object can't change with scale; this is unlike the idea of scale in cartographic maps, where a larger scale usually implies in smaller object support. In particular, the scale space of an image will be a sequence of images, *all of the same size of the original image*. When the scale increases, the images get more blurred, as if we were watching the images from a larger distance, but the support's size is kept fixed. The blurring is a direct consequence of the causality principle, that forces the image to lose detail as the scale increases.

The structure of this paper goes as follows:

In section 2 we will formally define (non-discrete) linear (Gaussian) scale-spaces and present their main properties.

In section 3, we will discretize such linear scale-spaces by discretizing the Gaussian filter kernel.

In section 4, we discretize linear scale-spaces by discretizing the PDE associated with the Heat Equaton, always keeping an eye to see if the non-discrete properties transfer to the discrete case. In particular, we present a new method (Crossed Convolutions) that solves the discrete version of the heat equation exactly, and therefore preserves many of the desirable properties of the non-discrete case.

Finally, in section 5, we compare the results of our implementations of the different methods (looking at theoretical differences, running time and output differences of the many scale-spaces).

1530-1834/01 $10.00 © 2001 IEEE

2 Gaussian Scale-Spaces

Definition 1 *Let $f : \mathbb{R}^n \to \mathbb{R}$ The **Gaussian scale-space** of f is the function $L : \mathbb{R}^n \times \mathbb{R}^+ \to \mathbb{R}$ given by*

$$L(\mathbf{x}, t) = f * G_t(\mathbf{x})$$

*where $L(\mathbf{x}, 0) = f(\mathbf{x})$, $G_t(\mathbf{x}) = \frac{1}{(4\pi t)^{\frac{n}{2}}} e^{-\frac{1}{4t}(x_1^2 + \ldots + x_n^2)}$ is the n-dimensional Gaussian distribution with variance $2t$ (standard deviation $\sigma = \sqrt{2t}$) and $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$. The parameter t is named **scale**.*

Figure 1 shows samples of Lenna's image scale-space at scales $t = 0, 0.25, 0.5, 1, 2, 4, 8, 16, 32$ (note that $t = 0$ corresponds to the original image).

Figure 1: Lenna's image scale-space.

The Gaussian scale-space has several interesting properties, most of them inherited from the Gaussian function. Let's list them briefly.

2.1 Properties

2.1.1 Linearity and Translation Invariance

Since we defined the Gaussian scale-space from a convolution, its dependence on f is necessarily linear and translation invariant.

2.1.2 Isotropy

Since the n-dimensional Gaussian is rotationally invariant or isotropic (see figure 2), so is the Gaussian scale-space.

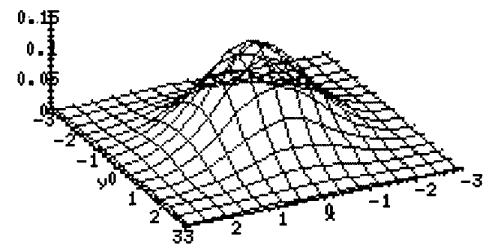

Figure 2: 2D Gaussian, $t = 0.5$.

2.1.3 Semigroup

The Gaussian kernels G_t have a semigroup property, that is, $G_{t_1+t_2} = G_{t_1} * G_{t_2}$ for any t_1 and $t_2 \in \mathbb{R}_+$. Therefore, $L(x, t_1 + t_2) = L(x, t_1) * G_{t_2 - t_1}(x)$, indicating *scale homogeneity*, that is, all scales of L are treated in the same way.

2.1.4 Causality

The causality property requires that the scale-space be somehow smoothed with increasing scale, but it can be formalized in many ways. One way of defining causality is to use a **maximum principle**: broadly speaking, the value of a local maximum (for fixed t) of $L(\cdot, t)$ does not increase as the scale t increases, and the value of a local minimum of $L(\cdot, t)$ does not decrease as the scale t increases. With such definition, the Gaussian scale-space is indeed *causal*.

2.1.5 Heat Equation

The Gaussian scale-space $L(\mathbf{x}, t) = f * G_t(\mathbf{x})$ is the solution of the following partial differential equation, known as the **Heat Equation**

$$\begin{cases} L(\mathbf{x}, 0) = f(\mathbf{x}) \\ \frac{\partial L(\mathbf{x}, t)}{\partial t} = \nabla^2 L(\mathbf{x}, t) \end{cases}$$

where the Laplacian on the right side is taken only with respect to the $\mathbf{x}$ variables ($\mathbf{x} \in \mathbb{R}^n$).

2.1.6 Uniqueness

The Gaussian scale-space is the only non-trivial space that is linear, isotropic, invariant by translations and satisfies the maximum principle stated above. For a proof, see [2].

3 Discretization via Convolution

How to discretize the Gaussian scale-space and keep the properties we listed in the previous section? As a first approach, we could define a discrete scale-space of a signal by applying any family of filters to it.

Definition 2 *Let g_t be a family of discrete filter kernels (with $g_0[n] = \delta[n] = \delta_{0n}$, the discrete Dirac function). The **discrete scale-space** (according to g_t) of a discrete signal $f : \mathbb{Z} \to \mathbb{R}$ is the function $L : \mathbb{Z} \times \mathbb{R}^+ \to \mathbb{R}$ given by*

$$L[n,t] = f * g_t[n]$$

Note that $L[n,0] = f[n]$.

The definition for multidimensional signals is similar. For example, for images we have:

Definition 3 *Let g_t be a family of discrete filter kernels (with $g_0[m,n] = \delta[m,n]$). The **discrete scale-space** (according to g_t) of an image $f : \mathbb{Z}^2 \to \mathbb{R}$ is the function $L : \mathbb{Z}^2 \times \mathbb{R}^+ \to \mathbb{R}$ given by*

$$L[m,n,t] = f * g_t[m,n]$$

Note that $L[m,n,0] = f[m,n]$.

Of course, this has the potential to be useless if the family g_t is not at least continuous on t, not to mention that we would like g_t to have *something* to do with the non-discrete Gaussian function! In this section, we present 3 alternatives for the construction of the family g_t, all of them trying to represent Gaussian scale-spaces (if possible, keeping the properties described in the continuous case).

3.1 Sampled Gaussian

Maybe the most straightforward solution to discretize the Gaussian scale-space is to take the family g_t from a direct uniform sampling of the Gaussian function. For example, in the unidimensional case, we could just calculate the value of the Gaussian function at the integers

$$g_t[n] = \frac{1}{\sqrt{4\pi t}} e^{-\frac{n^2}{4t}}$$

For implementation purposes, since this function has infinite support, we might have to decide on a cut-off point for n. It is customary to take this cut-off point at 3 or 4 standard deviations from the mean, that is, to take

$$g_t[n] = \begin{cases} \frac{C}{\sqrt{4\pi t}} e^{-\frac{n^2}{4t}}, \text{ for } |n| < 4\sqrt{\sigma} = 4\sqrt{2t} \\ 0, \text{ otherwise} \end{cases}$$

where the constant C is taken to make g_t normalized, that is, such that

$$\sum_{n=-\infty}^{\infty} g_t[n] = 1$$

This renormalization is specially relevant for small values of t, when the filter has small support. In particular, note that, when $t \to 0$, $G_t(0) \to \infty$ and we don't want the same behavior for $g_t[0]$. With the renormalization, $g_t[0] = 1$ for $t < 1/32$ (in fact, $g_t[n] = \delta[n]$ for $t < 1/32$).

Now we can define the scale-space of a 1D signal f as in the previous section[1]

$$L_t[n] = f[n] * g_t[n]$$

where $g_t[\cdot]$ is defined as above. Unfortunately, such scale-space does not have the semi-group nor the causality property (both properties are defined in the discrete case following the corresponding non-discrete definitions).

We extend the sampled Gaussian scale-space to the multi-dimensional case using tensor products, that is, using

$$g_t[n_1, n_2, ...] = g_t[n_1] . g_t[n_2] ....$$

where each of the g_t on the right side is defined according to the previous discussion.

Figure 3 shows samples of this discrete scale-space at variances 0 (original image), 1, 10 and 100.

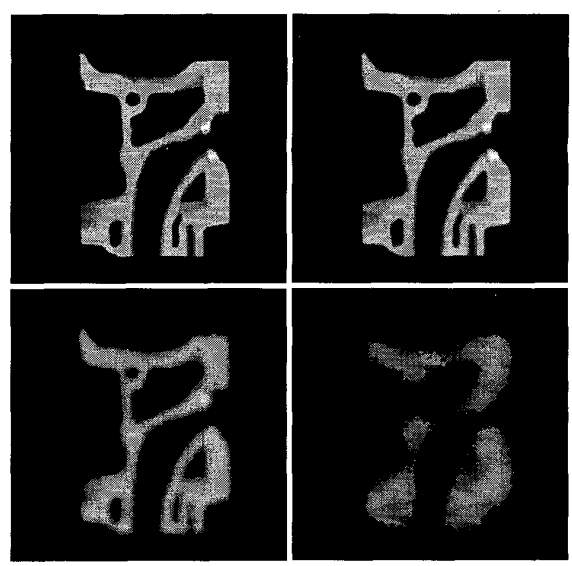

Figure 3: Sampled Gaussian scale-space.

[1] From here on, we slightly change the notation putting the variable t on $L[n,t]$ as a subscript.

3.2 Recursive Gaussian (Deriche)

Note that the calculation of the sampled Gaussian scale-space is very costly, specially when its variance is big, due to the computation of the convolutions involved. Deriche [4] shows how to approximate the sampled Gaussian kernel by a *recursive kernel* (that is, one whose convolution can be quickly calculated via recursive relations), namely,

$$\begin{aligned} e^{-\frac{1}{2\sigma^2}x^2} &\approx e^{-1.783\frac{x}{\sigma}}(1.68\cos(a\frac{x}{\sigma}) + 3.735\sin(a\frac{x}{\sigma})) - \\ &\quad e^{-1.723\frac{x}{\sigma}}(0.6803\cos(b\frac{x}{\sigma}) + 0.2598\sin(b\frac{x}{\sigma})) \\ &= g_{2t}(x) \end{aligned}$$

where $\sigma^2 = 2t$, $a = 0.6318$ and $b = 1.997$.

4 Discretization via Heat Equation

As we have seen in the previous section, the Gaussian scale-space can be defined as the solution of the heat equation $\frac{\partial L(\mathbf{x},t)}{\partial t} = \nabla^2 L(\mathbf{x}, t)$. This suggests a discrete scale-space based on the discretization of this equation instead of a direct discretization of the convolution in the solution. In fact:

Theorem 4 *The discrete heat equation*

$$\frac{\partial L_t[\mathbf{n}]}{\partial t} = \mathcal{A} * L_t[\mathbf{n}] \tag{1}$$

(where $\mathcal{A}$ is a multiple of the discrete Laplacian operator) defines the only scale-space that is linear, symmetric, translation-invariant, has the semigroup property and satisfies the maximum principle.

Proof. See chapter 10 of [2] or chapter 4 of [3]. Here we present just the general idea of the proof. First, we note that the semigroup property is basically equivalent to ask the scale-space to be given by some differential equation of the form

$$\frac{\partial L_t[\mathbf{n}]}{\partial t} = \mathcal{F}\{L[\cdot], \mathbf{n}\}$$

The invariance by translations practically eliminates the dependence on $\mathbf{n}$ on the right side; together with the linearity, we can already see that the equation must have a convolution form

$$\frac{\partial L_t[\mathbf{n}]}{\partial t} = \mathcal{A} * L_t[\mathbf{n}]$$

The maximum principle forces the operator $\mathcal{A}$ to have several properties; in particular, it has to be local (for example, in $\mathbb{Z}^2$, using 8-connectivity, $\mathcal{A}$ must be a $3x3$ kernel), its center coefficient is negative and all others are positive, and the sum of its coefficients is 0. Finally, the symmetry condition forces $\mathcal{A}$ to be symmetric as well. Summarizing, $\mathcal{A}$ must have the form

$$\mathcal{A} = K[1, -2, 1]$$

for one-dimensional scale-spaces, while

$$\mathcal{A}_c = K\begin{bmatrix} c & 1-2c & c \\ 1-2c & -4+4c & 1-2c \\ c & 1-2c & c \end{bmatrix} \qquad 0 \le c \le \frac{1}{2}$$

for the two-dimensional case. It is not hard to see that $\mathcal{A}$ is a multiple of a discrete approximation for the Laplacian operator. ■

Since the multiplicative constant K corresponds to a simple rescaling of the t parameter, it is common to take $K = \frac{1}{h^2}$ (where h is the distance between grid points) or, more simply, to take $K = 1$. The c parameter, on the other hand, cannot be determined without further constraints; note that c measures somewhat the "importance" of the diagonals of the matrix $\mathcal{A}$. It is common to take $c = 0$ for the simplest discretization of the Laplacian

$$\mathcal{A}_0 = \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

that has the advantage of satisfying the maximum principle even using 4-connectivity. However, for a more "isotropic" Laplacian, we prefer $c = \frac{1}{6}$:

$$\mathcal{A}_{\frac{1}{6}} = \frac{1}{6}\begin{bmatrix} 1 & 4 & 1 \\ 4 & -20 & 4 \\ 1 & 4 & 1 \end{bmatrix}$$

Whatever the choice of $\mathcal{A}$ is, we call it the *infinitesimal generator* of the scale-space. Based on this continuous scale, discrete space equation, we can come up with some other ideas for discrete scale-spaces.

4.1 First-order generator

We could apply a first-order method in t to solve equation 1; for example, in the 2D case, we take a small step Δt and write

$$L_{t+\Delta t}[m, n] \approx (\delta + \mathcal{A}\Delta t) * L_t[m, n]$$

where δ is the discrete unitary impulse

$$\delta = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Lindeberg [3] suggests then a discrete scale space computed by powers of a *first-order generator* $\delta + \mathcal{A}\Delta t$:

$$L_{t_0+n\Delta t}[m, n] \approx \underbrace{(\delta + \mathcal{A}\Delta t) * \cdots * (\delta + \mathcal{A}\Delta t)}_{n \text{ times}} * L_{t_0}[m, n]$$

Using the discrete 2D Laplacian above, it can be shown (like in [3]) that

- $\delta + \mathcal{A}_c \Delta t$ is unimodal only for $c \leq \frac{1}{4}$;
- $\delta + \mathcal{A}_c \Delta t$ it is separable only when $c = \frac{\Delta t}{2}$;
- The causality (maximum principle) is satisfied only if $0 \leq c \leq \frac{1}{4}$.
- $\delta + \mathcal{A}_c \Delta t$ has maximum "rotational symmetry" when $c = \frac{1}{6}$ (see chapter 4 of [3]).

Figure 4 show samples of this scale-space (using $c = \frac{1}{6}$) for $t = 0, 0.5, 5$ and 50.

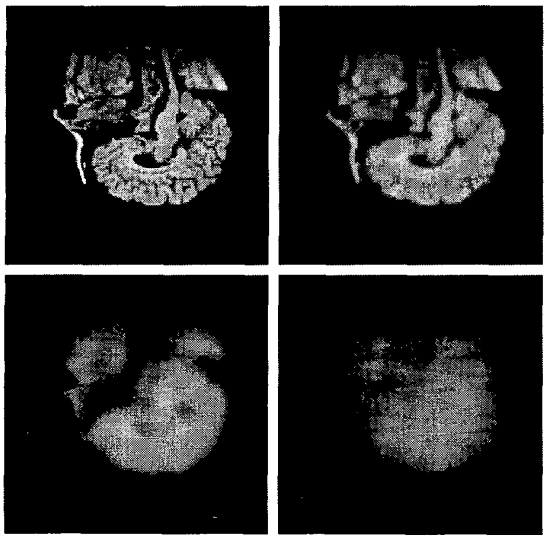

Figure 4: First-order generator scale-space

4.2 Crossed Convolutions

While all the scale-spaces presented so far approximate the original Gaussian scale-space, none of them solve exactly equation 1, and therefore none of them have all the properties that such solution would have. Moreover, it is really not absolutely necessary to use first-order approximations for equation 1: one can solve it explicitly.

For example, in the 1D case,

$$\begin{aligned} \frac{\partial L_t[n]}{\partial t} &= [1, -2, 1] * L_t[n] \\ L_0[n] &= f[n] \end{aligned}$$

can be explicitly solved using a Z-transform (see [2] or [3]). The solution is given by

$$L_t[n] = f * P_t[n]$$

where $P_t[n]$ is a *symmetric Poisson kernel*[2], that is,

$$\begin{aligned} P_t[n] &= e^{-2t} I_n(2t) \\ &= e^{-2t}(\cdots, I_{-n}(2t), \cdots, I_0(2t), \cdots, I_n(2t), \cdots) \end{aligned}$$

where $I_n(t)$ is the modified Bessel function [3] of order n.

Note that this is the **only 1D discrete scale-space** with the properties of linearity, invariance by translations, semi-group, symmetry and causality (given by the maximum principle). Also, it is not hard to show that $P_t[n]$ has variance $2t$, the same variance of the Gaussian $G_t(x)$. For these (and other) reasons, it is common to call $P_t[n]$ the *discrete Gaussian distribution.*

For 2D signals, we could just use tensor products; then, we would have a scale-space defined by

$$\begin{aligned} L_t[m,n] &= f * P_t[m,n] \\ P_t[m,n] &= P_t[m] P_t[n] \end{aligned}$$

that is indeed the solution of

$$\frac{\partial L[m,n]}{\partial t} = \mathcal{A}_0 * L[m,n]$$

However, can we write an explicit formula for the solution of the discrete heat equation for other discretizations of the Laplacian? The answer is positive:

Theorem 5 *The solution of the discrete 2D heat equation*

$$\begin{aligned} \frac{\partial L_t[m,n]}{\partial t} &= \mathcal{A}_c * L_t[m,n] \\ L_0[m,n] &= f[m,n] \end{aligned} \tag{2}$$

can be written as

$$L_t[m,n] = P^x_{(1-2c)t} * P^y_{(1-2c)t} * P^{x+y}_{ct} * P^{x-y}_{ct} * f[m,n]$$

where

$$\begin{aligned} P^x_\alpha[m,n] &= P_\alpha[m]\,\delta[n]; \\ P^y_\alpha[m,n] &= P_\alpha[n]\,\delta[m] \\ P^{x+y}_\alpha[m,n] &= P_\alpha[m]\,\delta[m-n]; \\ P^{x-y}_\alpha[m,n] &= P_\alpha[m]\,\delta[m+n] \end{aligned}$$

For example, P^x_α is the symmetric Poisson kernel of variance 2α spread only in the "x direction", while P^{x+y}_α is the same symmetric Poisson kernel but this time spread only in the direction $x = y$ (the 45^0 diagonal), with zeroes elsewhere.

[2] So called because it corresponds to the difference of two independent random variables with the same Poisson distribution.

[3] One convenient way to define the modified Bessel function is to look at the coefficients of the expansion of $e^{\frac{\alpha}{2}(z+z^{-1})}$ in a Laurent power series in z, that is

$$e^{\frac{\alpha}{2}(z+\frac{1}{z})} = \sum_{k=-\infty}^{\infty} I_n(\alpha) z^n$$

A subroutine for the calculation of modified Bessel functions can be found in [7].

Proof. Let

$$g_t(x,y) = \sum_{j,k=-\infty}^{\infty} L_t[j,k]\, x^j y^k$$

be the Z-transform of L_t. Then, from 2

$$\begin{aligned}\frac{\partial g_t(x,y)}{\partial t} &= t\{c\left(xy + xy^{-1} + x^{-1}y + x^{-1}y^{-1}\right) + \\ &\quad (1-2c)\left(x + x^{-1} + y + y^{-1}\right) - 4 + 4c\} \\ &\quad g_t(x,y)\end{aligned}$$

We solve this and get

$$g_t(x,y) = w_t(x,y) g_0(x,y)$$

where $w_t(x,y)$ is

$$e^{t\{c(xy+xy^{-1}+x^{-1}y+x^{-1}y^{-1})+(1-2c)(x+x^{-1}+y+y^{-1})-4+4c\}}$$

Rearranging the terms of $w_t(x,y)$ conveniently:

$$\begin{aligned}g_t(x,y) &= \\ &= e^{(1-2c)t(x+\frac{1}{x}-2)} \\ &\quad e^{(1-2c)t(y+\frac{1}{y}-2)} \\ &\quad e^{ct(xy+\frac{1}{xy}-2)} \\ &\quad e^{ct(\frac{x}{y}+\frac{y}{x}-2)} g_0(x,y)\end{aligned}$$

Now it is easy to see that each of the 4 exponentials is the Z-transform of one of the symmetric Poisson kernels in the solution (just expand them in Laurent power series and use the definition of the modified Bessel functions). ■

Therefore, we can compute $L_t[m,n]$ in 4 steps:

- Convolve $f[m,n]$ with $P^x_{(1-2c)t}$; this is the same as convolving each line of f with the 1D Poisson kernel $P_{(1-2c)t}$;
- Convolve each column of the result with $P_{(1-2c)t}$;
- Convolve each 45^0 diagonal of the result with P_{ct};
- Convolve each -45^0 diagonal of the result with P_{ct}.

In particular, the case $c = 0$ gives the old tensor product $L_t[m,n] = f[m,n] * P^x_t[m] * P^y_t[n]$.

Figure 5 shows samples of the Crossed convolution Scale-space for $c = \frac{1}{6}$ in the variances $2t = 0, 1, 10$ and 100. Since the convolution has infinite support, we mirrored the image across its borders to compute it.

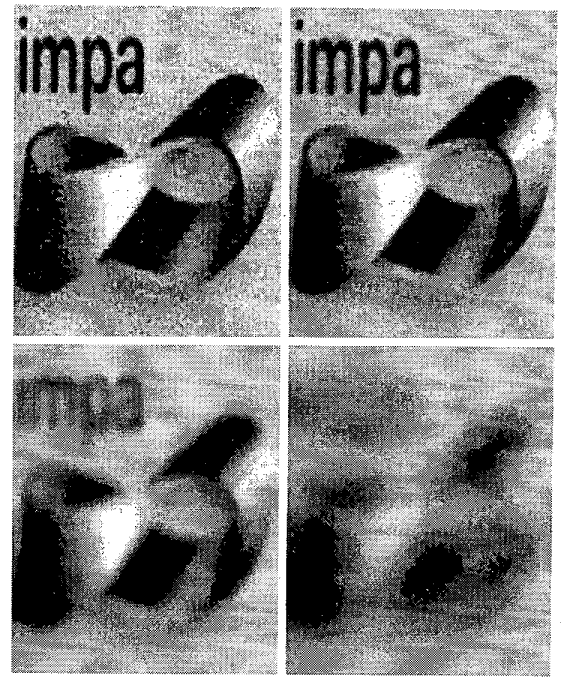

Figure 5: Crossed Convolution Scale-space

5 Comparative Analysis

In this section we make a comparative analysis of the discretizations.

5.1 Theoretical Analysis

The table below summarizes the properties of the considered scale-spaces

	semigroup	causality	scale
sampled Gaussian			$\mathbb{R}$
first-order generator[4]	✓	✓	$\Delta t\mathbb{Z}$
crossed convolutions	✓	✓	$\mathbb{R}$

All these methods try to approximate the continuous Gaussian scale-space in different ways. Note that both the sampled gaussian method and the crossed convolutions method incur in additional numerical errors by performing a kernel cut-off[5] to avoid infinite supports; the first-order generator method trades this infinite support by the approximation $L_{t+\Delta t}[m,n] \approx (\delta + \mathcal{A}\Delta t) * L_t[m,n]$.

[5] It is interesting to note that the cut-off for the crossed convolution method can be avoided using an implementation in the frequency domain, since the Fourier transforms of the symmetric Poisson kernels can be explicitly written.

5.2 Runtime

In the table below we show the algorithmic time and the runtime (in seconds) of the calculation of each scale-space for the indicated variances. While these values were taken from a low-performance computer, they are a good basis of comparison between different algorithms. All of them show that calculation of scale-spaces can be quite costly.

	mult	var=1.0	var=243.0
s. Gaussian	$8\sigma N$	0.22	1.87
first-order gen.	$10\sigma^2 N$	0.4	67.6
crossed conv.	$16\sigma N$	$0.22 - 0.36$	$1.87 - 2.37$

The first column indicates the number of multiplications for the calculation of each convolution, where N is the number of pixels in the image, σ is the standard deviation of the filter kernel and d is the denominator of the scale $\frac{n}{d}$ in the spline case.

For the crossed convolutions, we display running times for $c = 0$ and $c = \frac{1}{6}$. Note that, for $c = 0$, the times are the same as the sampled Gaussian times (since there are only two convolutions to be performed in this case).

Note also that, since the first-order generator method and the crossed convolutions method both have the semigroup property, we could calculate scale-spaces for larger scales using previously calculated smaller scales – this reduces running time for these algorithms when calculating several scales.

5.3 Output difference

We show the squared distance between images $f, g : \{0, 1, \cdots, M-1\} \times \{0, 1, \cdots, N-1\} \to \mathbb{R}$ given by

$$dist\,(f, g) = \frac{1}{MN} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} \left(f\,[m, n] - g\,[m, n]\right)^2$$

In the table below, we display the *dist* function for several variances and pairs of methods (the image used in these comparisons was Lenna's image). Signal extensions are by mirroring across borders unless stated otherwise.

		var	*dist*
S. Gaussian	C. Conv ($c = 0$)	$3,0$	$0,115$
S. Gaussian	C. Conv ($c = 0$)	$27,0$	$0,003$
C. Conv ($c = 0$)	1-Gen ($c = 0$)	$27,0$	$0,0044$
C. Conv ($c = 0$)	1-Gen ($c = \frac{1}{6}$)	$27,0$	$0,0029$
1-Gen ($c = \frac{1}{6}$)	C. Conv ($c = \frac{1}{6}$)	$27,0$	$5,074$

We note that all scale-spaces calculated are very close to each other. The only case in which two methods show some difference is the crossed convolutions method and the first-order generator method for $c = \frac{1}{6}$.

5.4 Final remarks

We note that the crossed convolutions method is the only one that retains all the properties of the continuous case. Since its computational cost is still acceptable, it is our personal favorite scale-space.

The sampled Gaussian method, while presenting similar results to the others and similar running times, seems to us to be overshadowed by the crossed convolutions method, that has the same running times and several theoretical properties in its favor.

First-order generators take longer and correspond only to a first-order approximation of the solution of the discrete heat equation. Besides, its theoretical properties of causality and semigroup breakdown whenever we need scales that are not multiples of the predefined scale step.

Finally, we must note that all the convolutions could be computed using Discrete Fourier Transforms (FFT algorithms); this might not only speed up all algorithms but, in some cases, avoid numerical errors that come from truncating the supports of the used kernels.

References

[1] Anderson Mayrink da Cunha, *Espaços de Escala e Detecção de Arestas*, M.Sc. Dissertation, IMPA, Rio de Janeiro, October 2000. http://www.visgraf.impa.br/escala.html

[2] Ralph Costa Teixeira, *Introdução aos Espaços de Escala*, Escola de Computação 2000, IME-USP, São Paulo, July 2000. http://www.visgraf.impa.br/Courses/eescala/index.html

[3] Tony Lindeberg, *Scale Space Theory in Computer Vision.* Kluwer Academic Publishers, Dordrecht, The Netherlands, 1994

[4] Rachid Deriche, *"Recursively implementing the gaussian and its derivatives,"* Tech. Report 1893, Programme 4 - Robotique, Image et Vision, INRIA - Institut National en Informatique et en Automatique, April 1993.

[5] Michael Unser, *"Splines: a perfect fit for signal and image processing,"* IEEE Signal Processing Magazine, vol. 16, no. 6, pp. 22-38, November 1999. http://bigwww.epfl.ch/publications/unser9902.html

[6] Yu-Ping Wang and S. L. Lee, *"Scale-Space Derived From B-Splines,"* IEEE Trans. Pattern Analysis and Machine Intelligence, vol. 20, no. 10, pp. 1040-1055, October 1998. http://wavelets.math.nus.edu.sg/~wyp/download_papers/Preprints.html

[7] William H. Press et al., *Numerical Recipes in C: the art of scientific computing.* second ed., Cambridge University Press, New York, 1992.