Modeling and Image Processing

Lionel Moisan

09/2012

Note: This document is a working and uncomplete version, subject to errors and changes. It was intended as course notes for 2nd year master students. Readers are welcome to point out mistakes to `lionel.moisan@parisdescartes.fr`*. Please do not make copies without this note.*

Note: ceci est un document de travail, régulièrement mis à jour et pouvant contenir des erreurs. Il a été originellement conçu comme support de cours à l'intention d'étudiants en 2ème année de master. Vous pouvez me signaler d'éventuelles erreurs, imprécisions, etc. en m'envoyant un courrier électronique à l'adresse `lionel.moisan@parisdescartes.fr`*. Prière de laisser ce message si vous reproduisez ce document.*

Contents

Chapter 1

Formation and representation of digital images

1.1 The pinhole camera model

1.1.1 Perspective projection

The first principle that governs the formation of an image in an optical device is geometry. In the classic **pinhole camera model**, the light rays coming from a point (X, Y, Z) of the scene are constrained to pass through the optical center $O = (0, 0, 0)$ and to hit the focal plane $Z = -f$ at point $(-fX/Z, -fY/Z, -f)$, producing in this plane an inverted image of the scene. In practice, it is more convenient (and geometrically equivalent) to assume that a non-inverted image is produced by the intersection of the light rays with the pseudo-focal plane $Z = f$ (see Figure 1.1).

Figure 1.1: The pinhole camera model

Definition 1 *The* **perspective projection** *of focal length f is the application*

$$\begin{array}{ccc}
\mathbb{R}^2 \times \mathbb{R}^* & \to & \mathbb{R}^2 \\
(X, Y, Z) & \mapsto & \left(f\dfrac{X}{Z}, f\dfrac{Y}{Z} \right).
\end{array}$$

The number f is called the focal length of the optics, and we shall call **geometrical image** the image $u(x, y)$ formed on the pseudo-focal plane, where $u(x, y)$ represent the density of energy received by the point (x, y, f) of this plane. Notice that u is assumed to be a monochromatic image here : to describe a color image, we should look of the distribution of this energy with respect to the wavelength and obtain color channels (typically, red-green-blue) by appropriate weightings of this distribution. We recall that the visible spectrum corresponds to wavelengths between 0.4μm (blue) and 0.8μm (red).

An important characteristic of the geometrical image formation process is that it is constrained by **occlusion** : when objects mask themselves each other, only the nearest is visible in a given direction. This has several consequences ; among them is the formation of edges (discontinuities) along the occlusion boundaries.

The pinhole camera model is an ideal model. In practice, some geometric distortions occur and physical 3D lines are not exactly projected into perfect 2D lines but are slightly curved, especially on the borders of the image. These distortions can be corrected by an appropriate post-processing of the image.

1.1.2 Projection of a plane

Consider a plane P of $\mathbb{R}^3$, endowed with an affine (or Euclidean) referential (M_0, U, V). We write $M_0 = (X_0, Y_0, Z_0)$, $U = (U_X, U_Y, U_Z)$, $V = (V_X, V_Y, V_Z)$ and we assume that $Z_0 > 0$. Any point $M(X, Y, Z) \in P$ can be written $M = M_0 + xU + yV$, that is

$$\begin{cases}
X = x\,U_X + y\,V_X + X_0 \\
Y = x\,U_Y + y\,V_Y + Y_0 \\
Z = x\,U_Z + y\,V_Z + Z_0
\end{cases}$$

If $Z \neq 0$, the perspective projection of this point is (x', y'), with

$$x' = f\frac{xU_X + yV_X + X_0}{xU_Z + yV_Z + Z_0}, \quad y' = f\frac{xU_Y + yV_Y + Y_0}{xU_Z + yV_Z + Z_0}. \tag{1.1}$$

By dividing by Z_0, this transform can be put under the general form below.

Proposition 1 *The perspective projection of a plane yields an* **homography***, that is a plane transform that can be written*

$$x' = \frac{ax + by + c}{dx + ey + 1}, \quad y' = \frac{a'x + b'y + c'}{dx + ey + 1}, \tag{1.2}$$

in affine or Euclidean referentials. Such a transform is defined by 8 parameters, or equivalently by the image of 4 points in generic position (3 of them are not aligned).

A special case is obtained when one assumes that the observed plane is "at infinity", that is when $Z_0 = +\infty$ (in practice, the approximation is correct when the variations of Z are small compared to Z_0 in the subdomain of P observed). Taking the limit $Z_0 = +\infty$ in (1.1), we obtain the **affine mapping**

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \frac{f}{Z_0} \begin{pmatrix} U_X & V_X \\ U_Y & V_Y \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \frac{f}{Z_0} \begin{pmatrix} X_0 \\ Y_0 \end{pmatrix}. \tag{1.3}$$

Such a plane transform depends on 6 parameters, and can be defined by the image of 3 non-aligned points. Notice also that any smooth plane transform T can be approximated locally at first order by an affine mapping, since a Taylor expansion yields

$$T(X) = T(X_0) + DT(X_0).(X - X_0) + \underset{X \to X_0}{O}(X - X_0),$$

where DT is the differential of T. For these reasons, affine mappings play an important role in image comparison and registration.

1.2 Point Spread Function

Aside from the geometric distortions mentioned above, there are another reasons why the image v formed on the focal plane differs from u. One of them is the fact that the image of a point source (a Dirac distribution of energy) will not yield a point, but a certain distribution of energy called the **Point Spread Function** (PSF), that is the impulse response of the optical device. The PSF arises from several phenomena that we are going to examine now. To this aim, we first need to recall some elements of Fourier theory.

1.2.1 Fourier Transform

In this section, we keep some generality by assuming that f is a real-valued function defined on $\mathbb{R}^d$ ($d = 2$ for images). As usual, L^1 is the space of functions f such that $\int_{\mathbb{R}^d} |f| < \infty$. The Fourier Transform of a function $f \in L^1$ is a continuous function defined by

$$\hat{f}(\xi) = \int_{\mathbb{R}^d} e^{-i \prec \xi, \mathbf{x} \succ} f(\mathbf{x}) \, d\mathbf{x}. \tag{1.4}$$

where $\prec \xi, \mathbf{x} \succ = \sum_{i=1}^{d} \xi_i x_i$ is the usual inner product. The function $\hat{f}$ is continuous and $\hat{f}(\mathbf{x}) \to 0$ as $|\mathbf{x}| \to \infty$. When $\hat{f} \in L^1$, the initial function f can be reconstructed from $\hat{f}$ using the inverse Fourier transform

$$f(\mathbf{x}) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} e^{i \prec \xi, \mathbf{x} \succ} \hat{f}(\xi) \, d\xi. \tag{1.5}$$

Note that this equation has to be taken in the L^1 sense, which means that f is almost everywhere equal to the continuous function defined by the right hand term.

Let us now assume that f belongs to the Schwartz Space $\mathcal{S}$, defined as the functions $f \in C^\infty$ such that $\mathbf{x}^\alpha \partial^\beta f(\mathbf{x}) \to 0$ as $|\mathbf{x}| \to \infty$ for any multi-indices $\alpha, \beta \in \mathbb{N}^d$. Then, it can be shown that $\hat{f}$ also belongs to $\mathcal{S}$, so that (1.5) still holds. The operator $\mathcal{F} : f \mapsto \hat{f}$ is an isomorphism of $\mathcal{S}$, and can be continuously extended into an isomorphism of L^2.

One great interest of the Fourier Transform is that it transforms convolutions in products and vice-versa. If f and g belong to $\mathcal{S}$, one has

$$\widehat{f \star g} = \hat{f} \cdot \hat{g} \qquad \text{and} \qquad \widehat{f \cdot g} = \frac{1}{(2\pi)^d} \hat{f} \star \hat{g}. \tag{1.6}$$

1.2.2 Diffraction

Since the aperture of the optical device is limited, it introduces diffraction. Its effect can be modeled by a convolution. In the following, we consider monochromatic incoherent light passing through a plane aperture. We assume that the aperture shape has a central symmetry and that the irradiation is constant on the aperture.

Theorem 1 (Fraunhofer Diffraction) *The physical image formed through a symmetric plane aperture is given by*

$$v = K_{diffraction} \star u,$$

where u is the (ideal) geometrical image, λ the wavelength, f the focal length and

$$K_{diffraction}(\mathbf{x}) = \frac{1}{2f^2} \left| \iint_{aperture} \exp\left(\frac{2i\pi <\mathbf{x}, \xi>}{\lambda f} \right) d\xi \right|^2. \tag{1.7}$$

In order to compute the diffraction kernel associated to a circular aperture (a disc with diameter D), we need to introduce the Bessel function of first kind and order $m \in \mathbb{Z}$, defined by

$$J_m(t) = \frac{1}{\pi} \int_0^\pi \cos(m\theta - t\sin\theta)\, d\theta, \tag{1.8}$$

and to recall the recurrence relation

$$\frac{d}{dt}\left(t^m J_m(t) \right) = t^m J_{m-1}(t). \tag{1.9}$$

Proposition 2 (circular aperture) *For a circular aperture with diameter D, the point spread function is*

$$K_{diffraction}(\mathbf{x}) = C \cdot \left(\frac{2J_1(r)}{r} \right)^2, \qquad \text{with} \quad r = \frac{\pi D |\mathbf{x}|}{\lambda f} \quad \text{and} \quad C = \frac{\pi^2 D^4}{32 f^2}. \tag{1.10}$$

Proof :

We apply (1.7) to the aperture domain $\{\xi \in \mathbb{R}^2,\ |\xi| \leqslant D/2\}$. Since $< R\mathbf{x}, R\xi >=< \mathbf{x}, \xi >$ for any rotation R, the PSF is radial, so that $K(\mathbf{x})$ only depends on $|\mathbf{x}|$. Writing $\xi = (\rho\cos\theta, \rho\sin\theta)$ and $\mathbf{x} = (x, y)$, we obtain

$$K(x,y) = \frac{1}{2f^2} \left| \int_0^{D/2} \int_0^{2\pi} \exp\left(\frac{2i\pi\rho(x\cos\theta + y\sin\theta)}{\lambda f} \right) \rho\, d\theta\, d\rho \right|^2,$$

and the previous symmetry remark leads us to

$$K(\mathbf{x}) = K(0, |\mathbf{x}|) = \frac{1}{2f^2} \left| 2\int_0^{D/2} \int_0^\pi \cos\left(\frac{2\pi\rho|\mathbf{x}|\sin\theta}{\lambda f} \right) \rho\, d\theta\, d\rho \right|^2.$$

Using the reduced variable r defined in (1.10), we then have, with $\rho' = 2r\rho/D$,

$$K(\mathbf{x}) = \frac{1}{2f^2} \left| \frac{2D^2}{4r^2} \int_0^r \int_0^\pi \cos(\rho'\sin\theta) \rho'\, d\theta\, d\rho' \right|^2 = \frac{\pi^2 D^4}{8f^2} \left| \frac{1}{r^2} \int_0^r J_0(\rho')\rho'\, d\rho' \right|^2.$$

We conclude by noticing that $\int_0^r J_0(\rho')\rho'\, d\rho' = rJ_1(r)$ thanks to (1.9). $\qquad \square$

The image $v = K_{diffraction}$, obtained for a point source, is called an Airy pattern, and is made of central bright spot surrounded by rings (see Figure 1.2). Since the first positive zero of $J_1(t)$ is located at $t \simeq 3.83 \simeq 1.22\pi$, the radius of the spot (Airy disc) is

$$r_a \simeq 1.22 \frac{\lambda f}{D}, \tag{1.11}$$

which gives an idea of the size of the smallest details that can be reproduced by an ideal optical device.

Example : let us compute an order of magnitude of the Airy spot. For a classical camera ($f/D = 5$) and in visible light ($\lambda = 0.6\mu m$), we obtain $r_a \simeq 3.6\mu m$, which is $1.5 \cdot 10^{-4}$ times the height of a full frame captor (24 mm).

(a)   (b)

Figure 1.2: (a) The diffraction kernel produced by a circular aperture (Airy pattern). The central disc receives 84% of the energy. (b) Its normalized radial profile, $r \mapsto \left(\frac{2J_1(r)}{r}\right)^2$.

1.2.3 Out of focus

For an optical device with non-zero aperture (the ideal pinhole camera model has zero aperture), the image of a point (X, Y, Z) is not formed in the plane $Z = f$ but on the plane $Z = t \geqslant f$, with t defined by Gauss' Formula

$$\frac{1}{t} = \frac{1}{f} - \frac{1}{Z}. \tag{1.12}$$

In general, the depth of the scene objects is not constant, so that it is impossible to obtain a perfect focus for all scene objects. The focusing operation consists in adjusting the position of the image plane in order to obtain a compromise. When an object with constant depth Z is de-focused in the plane $Z = t \pm \delta$, its image is blurred by a convolution with a PSF equal to the characteristic function of the aperture, scaled by a factor δ/f and normalized. In the general case, the image focused at $Z = t$ is

$$u_t(\mathbf{x}) = \frac{1}{S} \iint_{aperture} u\left(\mathbf{x} - \frac{|\delta(\mathbf{x})|}{f}\mathbf{y}\right) d\mathbf{y},$$

where

$$\frac{1}{t + \delta(\mathbf{x})} = \frac{1}{f} - \frac{1}{Z(\mathbf{x})},$$

S is the area of the aperture and $Z(\mathbf{x})$ is the depth of the point projected in $\mathbf{x}$ in the image plane. This transform is not a convolution in general. However, if all points of a scene should be ideally focused in $Z \geqslant t_0$, then for a circular aperture with diameter D one has

$$u_{t_0 - \delta} = K_{defocus} \star u_{t_0},$$

where

$$K_{defocus}(\mathbf{x}) = \frac{4f^2}{\pi \delta^2 D^2} \cdot \begin{cases} 1 & \text{if} \quad |x| \leqslant \dfrac{\delta}{f} \cdot \dfrac{D}{2}, \\ 0 & \text{else.} \end{cases} \tag{1.13}$$

Example : let us compute the defocus δ needed to produce a defocus kernel with the same radius as the Airy disc. According to (1.11) and (1.13), we have

$$1.22\frac{\lambda f}{D} = \frac{\delta D}{2f},$$

11

that is

$$\delta = 2.44\lambda \left(\frac{f}{D}\right)^2 \simeq 36\mu m$$

for $\lambda = 0.6\mu m$ (visible light) and $f/D = 5$. We can compare this number to the displacement of the focal plane required between focus at infinity and focus at, say, $Z = 1m$ for a focal of $f = 50mm$. We obtain, according to (1.12),

$$\delta = \frac{1}{\frac{1}{f} - \frac{1}{Z}} - f = \frac{f^2}{Z - f} \simeq 2600\mu m.$$

Thus, the precision required for focusing is, in this case, about 1% of the course of the focal plane.

1.3 Other distortions

In addition to the unavoidable diffraction, optical devices generally suffer from aberrations, mainly due to the fact that all light rays coming from a physical point do not converge in one point of the image plane. These aberrations are of several kinds

- **chromatic aberrations.** The focal f often depends on the wavelength, so that all color channels have not exactly the same focal plane, and for a given focus some channels will be slightly out of focus. This can be attenuated by using an achromatic doublet (a combination of two lenses that ensures the same focus for two given wavelengths) or by an appropriate post-processing of the image.

- **astigmatism and coma.** Some aberrations are non-uniform, in the sense that they depend on the image point. They result from the fact that all light rays coming from a given point of the scene do not converge to the same point of the image plane. These aberrations cannot be modeled with convolutions. Among them are astigmatism (different focal points for sagittal and tangential rays) and coma (different focal points for paraxial and marginal rays), which are responsible for non-isotropic impulse responses and for broader impulse responses on the borders of the image.

- **vignetting.** The density of light focused on the image plane is maximal on the center of the image plane (intersection with the optical axis), and slightly decreases as the distance from the center increases. This may cause an effect of "dark corners" on the resulting image, which is called vignetting.

Last, let us mention that long-focal devices (say $f \geqslant 1m$) suffer from **atmospheric turbulence**. In particular, ground-based astronomical observations are strongly limited by this phenomenon, and the order of magnitude of the maximum resolution allowed by atmospheric turbulence is about 1 arc second. To eliminate it, two solutions have been found in astronomy : spatial telescopes (e.g. Hubble) and adaptative optics (real-time physical deformation of the optical device).

1.4 Digitization

The image formed on the focal plane has to be recorded, either by a chemical or by an electronic device. Since we are interested in digital images, we shall only consider the second case. The digitization process is realized by a matrix of captors (in general, a CCD[1]) that covers the

---

[1]charge coupled device

focal plane. Each captor is an integrator : it counts the number of photons that hit the captor area during a fixed period. Ideally, the captors should cover completely the focal plane, and in practice this is often almost the case, at least for monochromatic images. The digitization of color images with only one focal plane is more challenging but we shall not go too far into these details here.

If C is the plane domain covered by a captor and δ the step of the grid, the recorded matrix is

$$I_{kl} = \iint_C u(k\delta + x, l\delta + y)\, dxdy.$$

Again, this can be modeled by a convolution : by using the PSF

$$K_{captor} = 1_C : x \mapsto \begin{cases} 1 & \text{if} \quad x \in C, \\ 0 & \text{else,} \end{cases}$$

we obtain

$$I_{kl} = (K_{captor} \star u)(k\delta, l\delta).$$

As before, this model is ideal and will suffer in practice from several distortions :

- **non-linearity.** The captor response is not exactly linear, especially for low and high energies. This means that instead of observing the matrix I_{kl}, we observe $g(I_{kl})$ for some non-linear increasing function g. A typical (normalized) function g is plotted on Figure 1.3.

Figure 1.3: The captor response is mostly linear except for extreme intensities.

- **noise.** Due to the random nature of the photo-electric effect (among other causes), fluctuations appear and the recorded image will be corrupted with (generally additive) noise.

- **quantization.** The recorded intensities have to be quantized on a finite number of levels, which produces a "quantization noise".

To sum up, a simple model for the image formation process is given by

$$I_{kl} = v(k\delta, l\delta), \qquad v = K \star u + n,$$

where

$$K = K_{captor} \star K_{diffraction} \star K_{defocus}$$

and $n(\mathbf{x})$ are random variables (generally i.i.d Gaussian). Notice that some other terms may have to be taken into account in the definition of K for low-quality optics. This model clearly shows two steps in the formation of digital images : the formation of v, defined on a continuous domain, and the sampling of v into the matrix I_{kl}. We are now going to study more precisely this sampling step.

1.5   Ideal sampling

The sampling process is usually described with Shannon Sampling Theorem, which says that a band-limited function can be exactly reconstructed from discrete samples taken on an appropriate grid. In the following, we shall need the Fourier Transform of a square window,

$$\mathcal{F}\left(1_{[-\pi,\pi]^2}\right)(\alpha,\beta) = 4\pi^2 \operatorname{sinc}(\alpha)\operatorname{sinc}(\beta), \tag{1.14}$$

with $\operatorname{sinc}(t) = \frac{\sin(\pi t)}{\pi t}$ for $t \neq 0$, and $\operatorname{sinc}(0) = 1$ by continuous extension. We also recall that if H_λ is the rescaling operator defined by

$$H_\lambda f(x,y) = f(\lambda x, \lambda y),$$

then one has

$$\mathcal{F} \circ H_\lambda = \frac{1}{\lambda^2} H_{1/\lambda} \circ \mathcal{F}.$$

1.5.1   Shannon Sampling Theorem

We first give a version of Shannon Sampling Theorem in the regular case ($f \in \mathcal{S}$) and for a critical sampling. We recall that $\operatorname{supp}(\hat{f})$ is the complement of the largest open set on which $\hat{f} = 0$.

Theorem 2 (Shannon) *Let $f \in \mathcal{S}$ and $\delta > 0$. If $\operatorname{supp}(\hat{f}) \subset [-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^2$, then*

$$\forall (x,y) \in \mathbb{R}^2, \qquad f(x,y) = \sum_{(k,l)\in\mathbb{Z}^2} f(k\delta, l\delta) \operatorname{sinc}\left(\frac{x}{\delta} - k\right) \operatorname{sinc}\left(\frac{y}{\delta} - l\right). \tag{1.15}$$

Proof :

1) We first prove the theorem for $\delta = 1$. Since $\operatorname{supp}(\hat{f}) \subset [\pi,\pi]^2$, the Fourier inversion formula writes

$$f(x,y) = \frac{1}{(2\pi)^2} \int_{[-\pi,\pi]^2} e^{i(\alpha x + \beta y)} \hat{f}(\alpha,\beta)\, d\alpha d\beta. \tag{1.16}$$

Let F be the 2π-periodized of $\hat{f}$ (along both coordinates). Since $\hat{f}$ is C^∞ on $\mathbb{R}^2$, all its derivatives vanish on the boundary of $[-\pi,\pi]^2$, which proves that F is C^∞ too. Consequently, the Fourier series of F converge normally and we have

$$F(\alpha,\beta) = \sum_{(k,l)\in\mathbb{Z}^2} c_{k,l}\, e^{i(\alpha k + \beta l)},$$

with

$$c_{k,l} = \frac{1}{(2\pi)^2} \int_{[-\pi,\pi]^2} e^{-i(sk+tl)}\, \hat{f}(s,t)\, dsdt = f(-k,-l)$$

thanks to (1.16). We conclude with Fubini Theorem that

$$\begin{aligned}
f(x,y) &= \frac{1}{(2\pi)^2} \int_{[-\pi,\pi]^2} e^{i(\alpha x + \beta y)} F(\alpha,\beta)\, d\alpha d\beta \\
&= \sum_{(k,l)\in\mathbb{Z}^2} f(k,l) \frac{1}{(2\pi)^2} \int_{[-\pi,\pi]^2} e^{i(\alpha x + \beta y) + i(-\alpha k - \beta l)}\, d\alpha d\beta \\
&= \sum_{(k,l)\in\mathbb{Z}^2} f(k,l) \frac{1}{(2\pi)^2} \cdot \mathcal{F}\left(1_{[-\pi,\pi]^2}\right)(k-x, l-y) \\
&= \sum_{(k,l)\in\mathbb{Z}^2} f(k,l) \operatorname{sinc}(k-x)\operatorname{sinc}(l-y)
\end{aligned}$$

thanks to (1.14).

2) Now we come to the general case ($\delta > 0$). If we assume that $\mathrm{supp}(\hat{f}) \subset [-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^2$, then the rescaled function $g = H_\delta f$ satisfies $\mathrm{supp}(\hat{g}) \subset [-\pi, \pi]^2$, so that we have

$$\forall (x, y) \in \mathbb{R}^2, \qquad g(x, y) = \sum_{(k,l) \in \mathbb{Z}^2} g(k, l) \, \mathrm{sinc}(x - k) \mathrm{sinc}(y - l),$$

or equivalently

$$\forall (x, y) \in \mathbb{R}^2, \qquad f(\delta x, \delta y) = \sum_{(k,l) \in \mathbb{Z}^2} f(k\delta, l\delta) \, \mathrm{sinc}(x - k) \mathrm{sinc}(y - l),$$

from which (1.15) follows after replacing (x, y) by $(x/\delta, y/\delta)$. $\qquad\qquad\square$

The proof above can be summarized very nicely using the theory of distributions. We first introduce the Dirac distribution $\delta_{x,y} \in \mathcal{S}'$, defined by $< \delta_{x,y} \, , \, \varphi >= \varphi(x, y)$ for any test function $\varphi \in \mathcal{S}$. Note that

$$(\delta_{x,y} \star \varphi)(s, t) = \varphi(s - x, t - y).$$

Now we need to know that the Fourier Transform of the Dirac comb

$$\Pi_s = \sum_{(k,l) \in \mathbb{Z}^2} \delta_{(ks, ls)} \qquad \text{is} \qquad \widehat{\Pi_s} = \left(\frac{2\pi}{s}\right)^2 \Pi_{\frac{2\pi}{s}}. \tag{1.17}$$

It is beyond the scope of this chapter to justify this precisely (we would have to define properly the Fourier transform of a periodic distribution), but let us mention that (1.17) follows directly from Poisson summation formula. Now the condition $\mathrm{supp}(\hat{f}) \subset [-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^2$, which is still true for the inverse Fourier transform of f, can be written

$$\mathcal{F}^{-1}(f) = (\mathcal{F}^{-1}(f) \star \Pi_{2\pi/\delta}) \cdot F_\delta, \qquad \text{with} \qquad F_\delta = 1_{[-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^2}. \tag{1.18}$$

Taking Fourier Transform on both sides yields, according to (1.6),

$$f = \frac{1}{(2\pi)^2} (f \cdot \widehat{\Pi_{2\pi/\delta}}) \star \widehat{F_\delta},$$

which boils down to

$$f = (f \cdot \Pi_\delta) \star S_{1/\delta} \qquad \text{with} \qquad S_{1/\delta}(\alpha, \beta) = \mathrm{sinc}\left(\frac{\alpha}{\delta}\right) \mathrm{sinc}\left(\frac{\beta}{\delta}\right). \tag{1.19}$$

Let us interpret Shannon Theorem. It basically says that a band-limited function can be reconstructed exactly from a discrete (but infinite) set of samples. The band-limited assumption is reasonable for most physical systems, that strongly attenuate frequencies outside a certain frequency domain. However, we could argue that the theorem does not tell us what happens when the energy is small (but non-zero) outside the frequency domain $[-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^2$. Moreover, the band-limited assumption is more questionable than in the case of signals (speech signals for example), since it is in contradiction with the fact that images have discontinuities at the boundaries of occluding objects.

The second objection to Shannon Theorem is that it requires an infinite set of samples to reconstruct the signal, whereas in practice only a finite number of them is available. This is quite disturbing since the sinc function, used for interpolation, has a slow decay, so that a lack of coefficients should yield an important error on the interpolation. We shall investigate this issue

in the next chapter, but let us mention that the slow decay of the sinc function can be avoided, provided that the image is oversampled. If $\text{supp}(\hat{f}) \subset [-\frac{\pi}{\delta'}, \frac{\pi}{\delta'}]^2$ with $\delta' > \delta$, the function $1_{[-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^2}$ in (1.18) can be replaced by a smooth (C^∞) function h such that

$$
\begin{array}{lll}
h(\alpha, \beta) = 0 & \text{if} & (\alpha, \beta) \notin [-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^2, \\
h(\alpha, \beta) = 1 & \text{if} & (\alpha, \beta) \in [-\frac{\pi}{\delta'}, \frac{\pi}{\delta'}]^2, \\
0 < h(\alpha, \beta) < 1 & \text{else.} &
\end{array}
$$

Then, (1.19) becomes

$$
f = (f \cdot \Pi_\delta) \star \left(\frac{\delta}{2\pi} \right)^2 \hat{h},
$$

and since h is smooth ($h \in \mathcal{S}$), its Fourier Transform will vanish faster at infinity than the sinc function ($\hat{h} \in \mathcal{S}$).

1.6 Aliasing

1.6.1 Spectrum periodization

A consequence of (1.17) is that sampling and periodization are dual operations with respect to Fourier Transform. Indeed, one has

$$
f \cdot \Pi_\delta \text{ (spatial sampling)} \quad \xrightarrow{\mathcal{F}} \quad \frac{1}{\delta^2} \cdot \hat{f} \star \Pi_{2\pi/\delta} \text{ (spectrum periodization)}
$$

which means that sampling an image is exactly equivalent to periodizing its spectrum. Now we can understand what happens if the condition $\text{supp}(\hat{f}) \subset [-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^2$ is not satisfied. The function

$$
g(x, y) = \sum_{(k,l) \in \mathbb{Z}^2} f(k\delta, l\delta) \operatorname{sinc}\left(\frac{x}{\delta} - k \right) \operatorname{sinc}\left(\frac{y}{\delta} - l \right)
$$

will still be equal to f for integer values of x/δ and y/δ, but will in general differ from f at other points. Precisely, the Fourier transform of g is obtained by adding to $\hat{f}$ its aliases obtained by the $\frac{2\pi}{\delta}$-periodization :

$$
\forall (\alpha, \beta) \in [-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^2, \qquad \hat{g}(\alpha, \beta) = \sum_{k,l \in \mathbb{Z}} \hat{f}(\alpha + 2k\pi/\delta, \beta + 2l\pi/\delta).
$$

This creates a phenomenon called **aliasing** : the spectrum of f will interfere with itself and some high-frequency components will be aliased to low frequencies (see Figure 1.4).

1.6.2 Case of a sine wave

To understand some effects of the aliasing on images, let us consider the particular case of a pure sine wave,

$$
u(\mathbf{x}) = \sin \prec \mathbf{k}, \mathbf{x} \succ .
$$

One easily checks that the Fourier Transform of u is

$$
\hat{u} = \frac{4\pi^2}{2i} \left(\delta_{\mathbf{k}} - \delta_{-\mathbf{k}} \right).
$$

Hence, u will be properly sampled on a grid with step δ as soon as $\mathbf{k} \in]-\frac{\pi}{\delta}, \frac{\pi}{\delta}[^2$. Let us now suppose that

$$
\mathbf{k} = \frac{2\pi}{\delta} \mathbf{n} + \mathbf{k}',
$$

well sampled image

aliased image (undersampled)
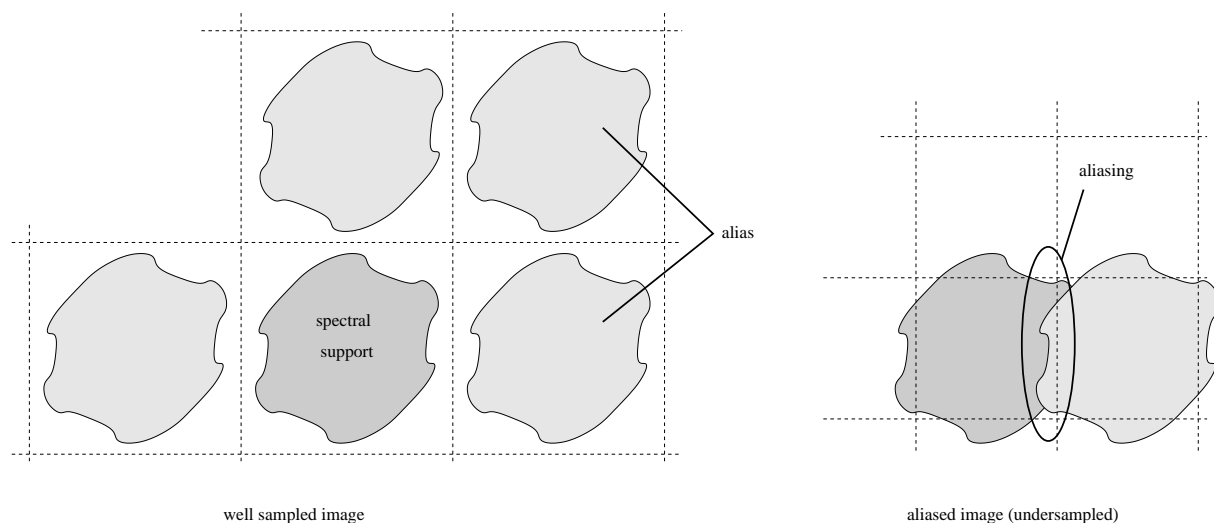
Figure 1.4: The aliasing process

with $\mathbf{n} \in \mathbb{Z}^2$ and $\mathbf{k}' \in]-\frac{\pi}{\delta}, \frac{\pi}{\delta}[^2$. Then the frequency $\mathbf{k}$ will be "aliased" into $\mathbf{k}'$, so that the image reconstructed from the discrete samples would be

$$\tilde{u}(\mathbf{x}) = \sin \prec \mathbf{k}', \mathbf{x} \succ .$$

This can create artificial patterns, especially for textures, when a high frequency $\mathbf{k}$ is aliased into a low frequency $\mathbf{k}'$.

This phenomenon can be very confusing in the case of movies, because the motion of such a texture will be distorted and amplified. The previous pattern moving (orthogonally to the stripes) at speed p,

$$u(\mathbf{x}, t) = \sin \left(\prec \mathbf{k}, \mathbf{x} \succ -tp|\mathbf{k}| \right),$$

will be aliased into

$$\tilde{u}(\mathbf{x}, t) = \sin \left(\prec \mathbf{k}', \mathbf{x} \succ -tp|\mathbf{k}| \right)$$

which is a sine wave moving at a speed

$$p' = p \cdot \frac{|\mathbf{k}|}{|\mathbf{k}'|}$$

that can be much larger than p.

1.7 Ringing

Shannon Theorem tells us how to avoid aliasing : ideally we should perform a perfect frequency cutoff, by transforming the geometrical image u into an image v such that

$$\hat{v} = \hat{u} \cdot 1_{[-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^2}. \tag{1.20}$$

Not only this perfect low-pass filter is impossible to realize physically in practice, but even if it was it would not be a good idea. Indeed, taking the inverse Fourier Transform of (1.20) yields

$$v = u \star R, \qquad \text{with} \qquad R(x, y) = C \cdot \text{sinc} \left(\frac{x}{\delta} \right) \text{sinc} \left(\frac{y}{\delta} \right)$$

for some constant C. This means that the hard frequency cutoff is equivalent in spatial domain to a convolution with a sinc function. Now, for a natural image containing edges, this will create a phenomenon called **ringing** : the edges will be replicated into vanishing parallel edges spaced by 2δ.

1.8 Realistic sampling

To avoid aliasing without creating too much ringing, a solution is to realize a smooth frequency cutoff, that is to sample the image

$$\hat{v} = \hat{u} \cdot \hat{h}, \tag{1.21}$$

where h is a smooth function such that $\hat{h} \simeq 0$ outside $[-\frac{\pi}{\delta}, \frac{\pi}{\delta}]^2$, and $\hat{h} \simeq 1$ inside. Since this process should occur before the digitization, it is quite difficult to design such a function h because of the physical constraints. As we saw before, for a perfect optical device we have

$$\hat{v} = \hat{u} \cdot \hat{K},$$

where

$$\hat{K} = \hat{K}_{captor} \cdot \hat{K}_{diffraction} \cdot \hat{K}_{defocus}.$$

Does this function satisfy the properties we expect ?

1.8.1 The diffraction term

For a circular aperture with diameter D, the diffraction term can be obtained by taking the Fourier Transform of (1.10), which gives

$$\hat{K}_{diffraction}(\xi) = \frac{\lambda^2 D^2}{4} \left(\arccos(\rho) - \rho\sqrt{1-\rho^2} \right) \qquad \text{with} \quad \rho = |\xi| \cdot \frac{\lambda f}{2\pi D}.$$

This function is plotted on Figure 1.5. The cutoff frequency $\rho_c = \frac{2\pi D}{\lambda f}$ corresponds to a sampling

Figure 1.5: The normalized $\hat{K}_{diffraction}$ term

step equal to $\delta_c = \pi/\rho_c = \frac{\lambda f}{2D}$, which is 0.41 (1/2.44) times the radius of the Airy spot,

$$r_a \simeq 1.22 \frac{\lambda f}{D}.$$

Let us compute some practical estimates.

- For a classical digital camera, the typical minimum aperture (largest f/D) is $f/D = 8$, which gives $\delta_c = 2.4\mu$ for visible light ($\lambda \simeq 0.6\mu$). A matrix of $R = 4N \times 3N$ square pixels with side δ_c will have a diagonal of $5N\delta_c$. For a 1/3 inch (8.5mm diagonal) CCD captor, the diffraction limit corresponds to the sampling limit when the resolution of the CCD is approximately $R = 6$ megapixels ($N = \frac{8500}{5 \times 2.4} \simeq 700$, $12N^2 \simeq 6.10^6$), which is the order of magnitude of digital cameras nowadays (note: this resolution should be roughly divided by 3 for color CCD, made of red, green and blue captors). A color digital camera with a larger resolution needs a 2/3 inch CCD captor to avoid oversampling at minimum aperture.

- For a 24×36 (non-digital) camera, the maximum diffraction allowed ($f/D = 22$) yields an Airy spot with diameter $32\mu m$ on the 24mm×36mm film, which corresponds to a diameter of 1.6mm on a giant 1.8m×1.2m (×50) poster. Hence, it will be difficult to observe the Airy spot, even with a very good film.

- For a telescope with diameter $D = 150$mm, the maximum angular resolution ($\arctan(r_a/f)$) allowed by diffraction is about 1 second of arc ($\tan(1'') \simeq 1.22 \cdot \frac{0.0006}{150}$), which is approximately the maximal resolution allowed by atmospheric turbulence in average conditions.

1.8.2 The captor term

If the captor domain is a square with side c (notice that we necessarily have $c \leqslant \delta$), the captor term is, up to a multiplicative constant,

$$\hat{K}_{captor}(\alpha, \beta) = \mathrm{sinc}\left(\frac{c\alpha}{2\pi}\right) \mathrm{sinc}\left(\frac{c\beta}{2\pi}\right).$$

The modulation brought by K_{captor} is represented on Figure 1.6. The cutoff frequency (first

Figure 1.6: The normalized $|\hat{K}_{captor}|$ term along each axis ($|\mathrm{sinc}|$ function)

zero) along each axis is $2\pi/c$.

How does this cutoff frequency compare to the Shannon limit π/δ ? To attain the critical value π/δ, we would like to have $c = 2\delta$, but this is physically impossible : we have at most $c = \delta$, since the captors cannot overlap. A nice solution (called hypomode) has been proposed by a CNES researcher named Bernard Rougé. It consists in simulating $c = 2\delta$ from $c = \delta$ by averaging the measured values on 2×2 pixels squares. The resulting digital image is more blurred but less aliased, so that a deconvolution algorithm may be able to restore more efficiently the true spectrum.

Chapter 2

Interpolation and geometric transforms

In this chapter, we investigate the problem of image interpolation : how a finite set of samples $u(k, l)$ should be interpreted to define a continuous image $v(x, y)$? This is the inverse of the sampling problem that we investigated in the previous chapter, and it has strong connexions with it.

Image interpolation has several applications : it is required for resolution changes (image enlargement) and for geometric transforms (non-integer translations, rotations, etc.). Image reduction, though not exactly a problem of interpolation, can be investigated using the same kind of tools. In this chapter, we shall only consider interpolations methods that are **exact** ($v(k, l) = u(k, l)$ for integer k, l), **linear** and translation-invariant. Such an interpolation is necessarily a convolution : if $u(k, l)$ is the initial discrete image, extended to $\mathbb{Z}^2$ in some way, then v can be written

$$v(x, y) = \sum_{(k,l) \in \mathbb{Z}^2} u(k, l) \, \phi(x - k, y - l),$$

with $\phi(0, 0) = 1$ and $\phi(k, l) = 0$ for $(k, l) \in \mathbb{Z}^2 \setminus \{(0, 0)\}$. Unless otherwise specified, we shall only consider in the following **separable symmetric** interpolations, where the function ϕ can be written

$$\phi(x, y) = \varphi(x)\varphi(y).$$

For that reason, we shall deal with one-dimensional signals (instead of two-dimensional images) in the following. We shall make use the 1D Dirac Comb,

$$\Pi_s = \sum_{k \in \mathbb{Z}} \delta_{ks},$$

whose Fourier Transform is

$$\hat{\Pi}_s = \frac{2\pi}{s} \Pi_{\frac{2\pi}{s}}.$$

2.1   The Fourier Model

The Fourier interpolation relies on the assumption that the discrete signal $u(k), k = 0..N - 1$ is the sampling of a band-limited signal satisfying Shannon condition. If we want to apply Shannon Theorem, we need to "guess" a value for the unknown samples $u(k)$ for $k \notin \{0..N - 1\}$, and expect that it will not influence too much the interpolation $v(x)$ for x inside $[0, N - 1]$. To each guess corresponds an interpolation method.

2.1.1 Sinc interpolation

The simplest possibility consists in setting the unknown values to 0. In this case, we get the continuous interpolation formula

$$\forall x \in \mathbb{R}, \qquad v(x) = \sum_{k=0}^{N-1} u(k) \operatorname{sinc}(x - k) \tag{2.1}$$

However, the "null" assumption is often not satisfactory, because it gives a particular meaning to the value 0, which is not always relevant. In particular, it has a major drawback : it does not reproduce constant signals : for example, the signal $1, 1, 1, ...1$ will be interpolated with (2.1) by an oscillating function.

2.1.2 Discrete sinc interpolation

Another possibility, which is commonly used, is to assume that the signal u is N-periodic, that is $u(k + N) = u(k)$ for all $k \in \mathbb{Z}$. This model has the advantage that it can be described with an elegant formalism : the Discrete Fourier Transform.

Definition 2 *The Discrete Fourier Transform (DFT) of the discrete signal $u(k), k = 0..N - 1$ is the discrete signal defined by*

$$\forall p \in \mathbb{Z}, \qquad \hat{u}(p) = \sum_{k=0}^{N-1} e^{-\frac{2i\pi pk}{N}} u(k). \tag{2.2}$$

Fourier Transform is one of the fundamental tools of signal and image processing. It has the following properties :

- it is N-periodic ($\hat{u}(p + N) = \hat{u}(p)$);

- it is conjugate-symmetric ($\hat{u}(-p) = \hat{u}(p)^*$) because u is real-valued (z^* means the complex conjugate of z).

- it is, up to a scale factor, an isometry of $l^2(\mathbb{R}^N)$, that is $\|\hat{u}\|_2 = \sqrt{N}\|u\|_2$; the Inverse Discrete Fourier Transform of $\hat{u}$ is

$$u(k) = \frac{1}{N} \sum_{p=0}^{N-1} e^{\frac{2i\pi pk}{N}} \hat{u}(p).$$

- if we define the discrete convolution of two N-periodic signals u and v by

$$u \star v(l) = \sum_{k=0}^{N-1} u(k)v(l - k),$$

then one checks easily that

$$\widehat{u \star v} = \hat{u} \cdot \hat{v} \qquad \text{and} \qquad \widehat{u \cdot v} = \frac{1}{N} \hat{u} \star \hat{v}.$$

What is the meaning of these numbers $\hat{u}(p)$? We saw that Shannon Theorem comes from the fact that the sampling process is equivalent to a periodization in Fourier domain : hence a

discrete signal may describe exactly the spectrum of a band-limited continuous signal. Here we have a discrete periodized signal

$$U = \sum_{k=0}^{N-1} u(k) \sum_{n \in \mathbb{Z}} \delta_{k+nN}$$

whose Fourier spectrum is

$$
\begin{aligned}
\hat{U}(\xi) &= \sum_{k=0}^{N-1} u(k) \int_{\mathbb{R}} \sum_{n \in \mathbb{Z}} \delta_{k+nN}(x) e^{-ix\xi} \, dx \\
&= \sum_{k=0}^{N-1} u(k) \int_{\mathbb{R}} \sum_{n \in \mathbb{Z}} \delta_{nN}(x-k) e^{-ix\xi} \, dx \\
&= \sum_{k=0}^{N-1} u(k) e^{-ik\xi} \int_{\mathbb{R}} \sum_{n \in \mathbb{Z}} \delta_{nN}(x) e^{-ix\xi} \, dx \\
&= \sum_{k=0}^{N-1} u(k) e^{-ik\xi} \cdot \widehat{\Pi_N}(\xi) \\
&= \sum_{k=0}^{N-1} u(k) e^{-ik\xi} \cdot \frac{2\pi}{N} \Pi_{2\pi/N}(\xi) \\
&= \sum_{p \in \mathbb{Z}} \hat{u}(p) \cdot \frac{2\pi}{N} \delta_{2\pi p/N}(\xi).
\end{aligned}
$$

Note that these are equalities between distributions and not between numbers. It proves that the Fourier Transform of U is, like U, a discrete and N-periodic signal, that can be written as a sum of Dirac masses weighted by the coefficients of the Discrete Fourier Transform of u.

Now we would like to recover the original signal $v(x)$ from which U have been sampled. Since

$$U = v \cdot \Pi_1 \qquad \Rightarrow \qquad \hat{U} = \hat{v} \star \Pi_{2\pi},$$

we can suppose that no aliasing occurred during the sampling process and recover the Fourier transform of v as a period of the Fourier transform of U. Let us suppose that $\hat{v}$ is supported by the interval

$$I_a = \frac{2\pi}{N} \cdot \{a, a+1, ..., a+N-1\}$$

for some $a \in \mathbb{Z}$. Using the Inverse Fourier Transform, we obtain that U has been sampled from

$$v_a(x) = \frac{1}{2\pi} \int_{\mathbb{R}} \sum_{p=a}^{a+N-1} \hat{u}(p) \cdot \frac{2\pi}{N} \delta_{2\pi p/N}(\xi) e^{ix\xi} \, d\xi = \frac{1}{N} \sum_{p=a}^{a+N-1} \hat{u}(p) e^{2i\pi px/N}. \qquad (2.3)$$

If $x \in \mathbb{Z}$, we recognize the Inverse Discrete Fourier Transform of $\hat{u}$ (that does not depend on a), and we check that $v_a(x) = u(x)$. Now if $x \notin \mathbb{Z}$, we can rewrite (2.3) as

$$
\begin{aligned}
v_a(x) &= \frac{1}{N} \sum_{p=a}^{a+N-1} \sum_{k=0}^{N-1} u(k) \, e^{\frac{2i\pi p(x-k)}{N}} \\
&= \frac{1}{N} \sum_{k=0}^{N-1} u(k) \, e^{\frac{2i\pi a(x-k)}{N}} \cdot \frac{1 - e^{2i\pi(x-k)}}{1 - e^{\frac{2i\pi(x-k)}{N}}} \\
&= \sum_{k=0}^{N-1} u(k) \, e^{\frac{2i\pi}{N}(x-k)[a+\frac{N-1}{2}]} \cdot \frac{\sin \pi(x-k)}{N \sin \frac{\pi}{N}(x-k)}.
\end{aligned}
$$

In order $v_a(x)$ to be real, we should have $a = -\frac{N-1}{2}$.

- If N is odd, this is possible, and everything goes well. The Shannon condition is satisfied since $\operatorname{supp}(\widehat{v_a}) = \{-\frac{N-1}{2} .. \frac{N-1}{2}\} \subset]-\frac{N}{2}, \frac{N}{2}[$, so that the periodization of the spectrum will introduce no aliasing, and we get the reconstruction formula

$$v(x) = \sum_{k=0}^{N-1} u(k) \, \frac{\sin \pi(x-k)}{N \sin \frac{\pi}{N}(x-k)}.$$

- If N is even, we have got a problem, since $\frac{N-1}{2}$ is not an integer. Indeed, we cannot decide if the value $\hat{u}(\frac{N}{2})$ should be attached to frequency $\frac{N}{2}$ or $-\frac{N}{2}$. This ambiguous situation is demonstrated, for example, in the fact that the two functions $x \mapsto \cos(\pi x + \phi)$ and $x \mapsto \cos(\phi)\cos(\pi x)$ have the same samples. Hence, it is necessary that the discrete signal satisfies $\hat{u}(\frac{N}{2}) = 0$, and then the reconstruction is obtained (among other possibilities, since $\hat{u}(\frac{N}{2}) = 0$ introduces dependencies) by averaging the $a = -N/2$ and $a = -(N-2)/2$ choices, yielding

$$v(x) = \sum_{k=0}^{N-1} u(k) \, \frac{\sin \pi(x-k)}{N \tan \frac{\pi}{N}(x-k)}.$$

Let us summarize the discussion above. If a signal has been sampled according to Shannon Theorem, then the sampled signal must belong to the space $\mathcal{S}_N$ defined below.

Definition 3 *We call $\mathcal{S}_N$ the space of discrete signals $u(k), k = 0..N-1$ such that either N is odd or $\hat{u}(N/2) = 0$.*

When N is even, the condition $\hat{u}(N/2) = 0$ also writes

$$\sum_{k=0}^{N-1} (-1)^k u(k) = 0, \tag{2.4}$$

since the left term is nothing but $\hat{u}(N/2)$. The dimension of $\mathcal{S}_N$ is $N-1$, and the orthogonal projection from $\mathbb{R}^N$ to $\mathcal{S}_N$ is defined by

$$\begin{aligned} P_N : \mathbb{R}^N &\to \mathcal{S}_N \\ u &\mapsto \left(k \mapsto u(k) - \frac{(-1)^k}{N} \hat{u}\left(\frac{N}{2}\right) \right). \end{aligned}$$

Now, under the assumption that the initial signal is N-periodic, the Shannon reconstruction is, as we saw above, given by a discrete version of the sinc function.

Proposition 3 *The Shannon N-periodic interpolation of a discrete signal $u \in \mathcal{S}_N$ is equal to*

$$v(x) = \sum_{k=0}^{N-1} u(k) \operatorname{sincd}_N(x-k), \tag{2.5}$$

where the discrete cardinal sine of order N is the N-periodic function defined by

$$\forall t \in \mathbb{R} \setminus N\mathbb{Z}, \qquad \operatorname{sincd}_N(t) = \begin{cases} \dfrac{\sin \pi t}{N \sin \frac{\pi}{N} t} & \text{if } N \text{ is odd,} \\[3mm] \dfrac{\sin \pi t}{N \tan \frac{\pi}{N} t} & \text{if } N \text{ is even,} \end{cases}$$

and $\mathrm{sincd}_N(kN) = 1$ *for* $k \in \mathbb{Z}$. *It is equivalent to define* v *by*

$$v(x) = \frac{1}{N} \sum_{p=a}^{a+N-1} e^{\frac{2i\pi px}{N}} \hat{u}(p) \tag{2.6}$$

with $a = -\frac{N-1}{2}$ *(N odd) or* $a = -\frac{N}{2}$ *(N even).*

Remark : we have shown that in order to interpret the spectrum of a discrete signal as the periodization of the finitely supported spectrum of a band-limited signal, a possibility consists in selecting the only period of the spectrum that is symmetric with respect to 0. If the selected frequency domain has to be connected (as a subset of $\mathbb{Z}$*), this is the only choice that yields real values for the interpolate. In two dimensions, the choice may be broader if we do not impose the isotropy of the coordinates : any symmetric connected domain that defines a tiling of* $\mathbb{Z}^2$ *by periodization along both coordinates is valid. This possibility has been recently used in satellite imaging, in a situation where the two coordinates were not equivalent (because of the proper motion of the satellite).*

2.1.3 Symmetrization and the Cosine Transform

Since the periodization of a signal may introduce a jump in the discrete values, and consequently oscillations in the interpolate, it is often interesting to perform a symmetrization of the discrete signal $u(k), k = 0..N - 1$. The result is a new signal $\tilde{u}(k), k = 0..2N - 1$ obtained by setting $\tilde{u}(2N - 1 - k) = \tilde{u}(k) = u(k)$ for $0 \leqslant k \leqslant N - 1$. Note that $\tilde{u}$ necessarily belongs to $\mathcal{S}_N$, since it clearly satisfies (2.4). Its Fourier Transform is

$$\begin{aligned}
\hat{\tilde{u}}(p) &= \sum_{k=0}^{2N-1} e^{\frac{-2i\pi pk}{2N}} \tilde{u}(k) \\
&= \sum_{k=0}^{N-1} \left(e^{\frac{-2i\pi pk}{2N}} + e^{\frac{-2i\pi p(2N-1-k)}{2N}} \right) \tilde{u}(k) \\
&= 2e^{i\pi p/2N} \sum_{k=0}^{N-1} \cos\left(\frac{\pi p}{N}(k + \frac{1}{2}) \right) u(k).
\end{aligned}$$

Hence, if we define the **Discrete Cosine Transform** (DCT) of u by

$$\forall p \in \mathbb{Z}, \qquad \mathcal{C}_u(p) = \sum_{k=0}^{N-1} \cos\left(\frac{\pi p}{N}(k + \frac{1}{2}) \right) u(k), \tag{2.7}$$

then one simply has

$$\hat{\tilde{u}}(p) = 2e^{i\pi p/2N} \cdot \mathcal{C}_u(p). \tag{2.8}$$

Note that this implies that $|\hat{\tilde{u}}|$ is, like $\mathcal{C}_u$, even.

Exercise : show that a signal u *can be recovered from* $\mathcal{C}_u$ *using the inverse DCT, that is*

$$u(k) = \frac{1}{N} \left[\mathcal{C}_u(0) + 2 \sum_{p=1}^{N-1} \cos\left(\frac{\pi p}{N}(k + \frac{1}{2}) \right) \mathcal{C}_u(p) \right].$$

2.1.4 The Fast Fourier Transform (FFT)

Before we apply Fourier interpolation to geometric transforms, let us highlight a very interesting property of the Discrete Fourier Transform : there exists a fast algorithm to compute it.

The Discrete Fourier Transform consists in N terms defined each by a sum of N terms, so that its complexity is *a priori* $O(N^2)$. However, we know since Cooley and Tukey (1965) that the DFT can be computed recursively more quickly, especially when N is a power of two. This is the Fast Fourier Transform (FFT) algorithm.

Suppose that $N = 2^n$. We want to compute the values $P(w^p)$, with $w = e^{-2i\pi/N}$ and $p = 0..N - 1$, and P being a polynomial with degree $N - 1$, precisely

$$P(X) = \sum_{k=0}^{N-1} u(k)X^k.$$

Let us call $C(N)$ the number of operations required (additions plus multiplications). We can write

$$P(w^p) = P_0(w^{2p}) + w^p P_1(w^{2p}), \tag{2.9}$$

where

$$P_0(X) = \sum_{k=0}^{N/2-1} u(2k)X^k \qquad \text{and} \qquad P_1(X) = \sum_{k=0}^{N/2-1} u(2k+1)X^k.$$

We have translated the task of order N (evaluation of P for the N-th roots of unity) into two similar tasks of order $N/2$ (evaluation of P_0 and P_1 for the $N/2$-th roots of unity), plus N multiplications and N additions (we suppose that the value of w^p has been computed in advance). The consequence is that

$$C(N) = 2C\left(\frac{N}{2}\right) + 2N.$$

Since $C(1) = 0$, we easily deduce that $C(N) = 2N\log_2(N)$. Hence, the recursive algorithm suggested above computes the Fourier transform of a discrete signal with size N in $O(N\log N)$ operations. For a two-dimensional image with size $N \times N$, the separability implies that $O(N^2 \log N)$ (instead of N^4) are enough.

What happens if N is not a power of two ? Following the same idea as (2.9), we get, when $N = qr$,

$$P(w^p) = P_0(w^{qp}) + w^p P_1(w^{qp}) + \ldots + w^{(q-1)p} P_{q-1}(w^{qp}),$$

where the P_i's have degree $r - 1$. Hence,

$$C(qr) = 2qr(q-1) + qC(r),$$

so that if the prime factor decomposition of N is $N = \Pi_i p_i$, then the Fourier Transform of N can be computed in

$$C(N) = 2N\sum_i (p_i - 1)$$

operations. Consequently, the algorithm is more efficient when N does not contain large prime factors.

In most cases, we are not dealing in practice with periodic signals or images, but we periodize them in order to use the convenient Discrete Fourier representation. For that reason, the signal or image dimensions may often be extended or reduced in order to match a power of two, or at least an integer with small prime factors.

2.1.5 Zoom (zero-padding)

Fourier interpolation can be easily applied to signal (and image) enlargement (zoom). According to Shannon Theorem, oversampling a signal by a factor $\lambda = M/N$ could be realized with (2.5). However, this is not a good idea : first, because we saw that formulations in Fourier domain generally yield more efficient algorithms based on FFT, secondly because in this case the Fourier formulation is very simple. Indeed, according to (2.6) the magnified signal should be defined in Fourier domain by

$$\forall p \in \mathbb{Z}, \ |p| < M/2, \qquad \hat{u}_\lambda(p) = \left\{ \begin{array}{ll} \hat{u}(p) & \text{if} \quad |p| < N/2 \\ 0 & \text{else.} \end{array} \right.$$

This is why this method is called **zero-padding**. When N and M are power of two, it can be realized in $O(M \log M)$ operations.

2.1.6 Non-integer translation

We saw that Fourier interpolation may be viewed as a direct consequence of Shannon Sampling Theorem. But it may be justified in a completely different way : this is the optimal way to perform non-integer translations without loosing information, as stated by the following theorem.

Theorem 3 *There exists a unique family of linear operators* $(T_z)_{z \in \mathbb{R}} : \mathcal{S}_N \to \mathcal{S}_N$ *such that :*

(i) $z \mapsto T_z$ *is continuous,*

(ii) $\forall k, z \in \mathbb{Z}, \ T_z u(k) = u(k - z),$

(iii) $\forall w, z \in \mathbb{R}, \ T_{w+z} = T_w \circ T_z,$

(iv) $\lim\limits_{z \to 0} |z|^{-1} \| T_z - id \|_2$ *is minimal.*

It is defined by

$$T_z u(k) = \sum_{n=0}^{N-1} u(n) \operatorname{sincd}_N(k - z - n) \tag{2.10}$$

or equivalently, by

$$\widehat{T_z u}(p) = e^{\frac{-2i\pi pz}{N}} \hat{u}(p) \qquad \left(-\frac{N}{2} < p < \frac{N}{2} \right). \tag{2.11}$$

Proof :

    • Since each operator T_z is linear and translation-invariant, it is a convolution, as we are going to show now. If N is odd, we can define

$$\forall k \in \mathbb{Z}, \forall n \in \{0..N-1\}, \qquad e_n(k) = \left\{ \begin{array}{ll} 1 & \text{if} \quad k = n \mod N \\ 0 & \text{else.} \end{array} \right.$$

and set

$$\alpha_z(k, n) = \left(T_z e_n \right)(k).$$

Using (ii) and (iii), we get, for all $(k, n) \in \mathbb{Z} \times \{0..N-2\}$,

$$\alpha_z(k, n) = \left(T_z e_n \right)(k) = \left(T_{-1} T_z T_1 e_n \right)(k) = \left(T_{-1} T_z e_{n+1} \right)(k) = \left(T_z e_{n+1} \right)(k+1) = \alpha_z(k+1, n+1).$$

It follows that $\alpha_z(k, n) = \alpha_z(k - n, 0)$, so that

$$(T_z u)(k) = \left(T_z \sum_{n=0}^{N-1} u(n) e_n \right)(k) = \sum_{n=0}^{N-1} u(n)(T_z e_n)(k) = \sum_{n=0}^{N-1} u(n) \beta_z(k - n)$$

with $\beta_z(n) = \alpha_z(n, 0)$. If N is even, we can define

$$\forall k \in \mathbb{Z}, \forall n \in \{0..N - 1\}, \qquad f_n(k) = e_n(k) + \frac{(-1)^{n-k+1}}{N}.$$

We have

$$\sum_{k=0}^{N-1} (-1)^k f_n(k) = \sum_{k=0}^{N-1} (-1)^k e_n(k) + \frac{1}{N} \sum_{k=0}^{N-1} (-1)^{n+1} = (-1)^n + (-1)^{n+1} = 0,$$

which proves that f_n belongs to $\mathcal{S}_N$. Since $T_1 f_n = f_{n+1}$, we can set

$$\alpha_z'(k, n) = \left(T_z f_n \right)(k)$$

and show as above that $\alpha_z'(k, n) = \alpha_z'(k - n, 0)$. Now, since for $u \in \mathcal{S}_N$ we have

$$\left(\sum_{n=0}^{N-1} u(n) f_n \right)(k) = \sum_{n=0}^{N-1} u(n) e_n(k) + \frac{1}{N} \sum_{n=0}^{N-1} u(n)(-1)^{n-k+1} = u(k) + \frac{(-1)^{1-k}}{N} \hat{u}(N/2) = u(k),$$

we deduce that

$$(T_z u)(k) = \left(T_z \sum_{n=0}^{N-1} u(n) f_n \right)(k) = \sum_{n=0}^{N-1} u(n)(T_z f_n)(k) = \sum_{n=0}^{N-1} u(n) \beta_z(k - n)$$

with $\beta_z(n) = \alpha_z'(n, 0)$.

• Since T_z is a convolution, it is diagonalized by the Discrete Fourier Transform, that is

$$\widehat{T_z u}(p) = \hat{\beta}_z(p) \hat{u}(p). \tag{2.12}$$

Notice that β_z is N-periodic, so we can assume $\beta_z \in \mathcal{S}_N$ without loss of generality, since for even N the value of $\hat{\beta}_z(N/2)$ has no influence on T_z ($\hat{u}(N/2) = 0$).

• Now, property (iii) yields

$$\forall z, w \in \mathbb{R}, \qquad \hat{\beta}_{z+w}(p) = \hat{\beta}_z(p) \hat{\beta}_w(p),$$

and by continuity of $z \mapsto \hat{\beta}_z(p)$ (property (i)) we get that

$$\beta_z(p) = e^{\gamma(p)z}, \qquad -\frac{N}{2} < p < \frac{N}{2}$$

for some $\gamma(p) \in \mathbb{C}$. Since $\hat{\beta}_1(p) = e^{\frac{-2i\pi p}{N}}$, we have

$$\gamma(p) = -2i\pi \left(\frac{p}{N} + k(p) \right),$$

where $k(p) \in \mathbb{Z}$ and $k(-p) = -k(p)$ (the fact that $T_z u$ is real-valued implies that $\beta_z(-p) = \beta_z(p)^*$).

- Last, we compute

$$
\begin{aligned}
\|T_z - id\|_2^2 &= \sup_{u \in \mathcal{S}_N, \|u\|_2 = 1} \|T_z u - u\|_2^2 \\
&= \frac{1}{N} \sup_{u \in \mathcal{S}_N, \|u\|_2 = 1} \|\widehat{T_z u} - \hat{u}\|_2^2 \\
&= \frac{1}{N} \sup_{u \in \mathcal{S}_N, \|\hat{u}\|_2^2 = N} \sum_{p=0}^{N-1} |e^{-2i\pi\left(\frac{p}{N} + k(p)\right)z} - 1|^2 \cdot |\hat{u}(p)|^2 \\
&= 4 \max_{-\frac{N}{2} < p < \frac{N}{2}} \sin^2\left(\pi\left(\frac{p}{N} + k(p)\right)z\right) \\
&= 4\pi^2 z^2 \max_{-\frac{N}{2} < p < \frac{N}{2}} \left(\frac{p}{N} + k(p)\right)^2 + \underset{z \to 0}{o}(z^2).
\end{aligned}
$$

Hence,

$$
\lim_{z \to 0} |z|^{-1} \|T_z - id\|_2 = 2\pi \max_{-\frac{N}{2} < p < \frac{N}{2}} \left|\frac{p}{N} + k(p)\right|
$$

is minimal if and only if $k(p) = 0$ for all p. We conclude that

$$
\beta_z(p) = e^{-2i\pi p z / N}, \qquad -\frac{N}{2} < p < \frac{N}{2}
$$

as announced. Inversely, one easily checks that the operators (T_z) defined by (2.10) satisfy the required properties (i), (ii), (iii) and (iv). □

We just proved that the only minimal continuous semi-group extending the integer translations is given by the Fourier model. Its main interest is that non-integer translations are realized, like integer translation, without any loss of information : T_z is an isometry of $\mathcal{S}_N$ for all z. Notice that like for the zoom, it is computationally more efficient to realize the translation using the Fourier formulation (2.11). It yields an algorithm with complexity $O(N \log N)$ when N is a power of two.

2.1.7   Rotation

The rotation of an image is *a priori* incompatible with the discrete Fourier model, because if an image is periodized along a system of coordinates, in general the rotated image will not be periodic along the same system (except if the rotation angle is a multiple of $\pi/2$, which is a trivial case). However, it is possible to use the Fourier model to perform very accurate rotations.

The first solution is to use the interpolation formula (2.5) — or, to avoid periodization artifacts, (2.1) — and to sample the new image

$$
v(k, l) = u(k\cos\theta - l\sin\theta + \alpha, k\sin\theta + l\cos\theta + \beta)
$$

(the translation term (α, β) indirectly defines the rotation center). However, since there is no way to separate the variables, this would require $O(N^4)$ operations, which is generally too much.

Fortunately, there exists an efficient and elegant solution, that has been found independently by several authors in the 80's. The idea is to decompose the rotation into the product of three "shears",

$$
\begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} = \begin{pmatrix} 1 & -\tan\frac{\theta}{2} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \sin\theta & 1 \end{pmatrix} \begin{pmatrix} 1 & -\tan\frac{\theta}{2} \\ 0 & 1 \end{pmatrix}, \tag{2.13}
$$

and to use the Fourier model to implement the shears as "variable translations". Let u be an image with size $N \times N$. For $s \in \mathbb{R}$, the typical shear $\begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix}$, applied to u, yields the image

$$v(k,l) = u(k + sl, l). \tag{2.14}$$

If we set $u_l(k) = u(k,l)$ and $v_l(k) = v(k,l)$, then (2.14) can be viewed as a translation and computed by (2.11), that is

$$\widehat{v_l}(p) = e^{2i\pi slp/N} \widehat{u_l}(p), \qquad -\frac{N}{2} < p < \frac{N}{2}.$$

Hence, if N is a power of two, the rotation of an image using (2.13) can be performed in $O(N^2 \log N)$ operations. The only drawback of this method is that some parts of the original image are "scrambled" during this operation. However, this only affects the borders of the images, so that an appropriate pre-processing (for example, enlarging the domain of u with a spatial zero-padding) can solve this problem.

2.2   Other interpolation methods

The Fourier model is an interesting theoretical tool but it suffers from not being local, in the sense that the interpolate of an image around a point depend on all the image samples. This may be questionable from a physical point of view, and has the drawback that in all situations where the Fast Fourier Transform cannot be used (e.g. to apply an homography to an image), Fourier interpolation requires heavy computations. Another reason why the Fourier model is not completely adequate is that it relies on the strong hypothesis that the considered image was made band-limited before it was sampled. If this is not the case, which happens in practice, the Fourier interpolate will present unwanted ringing artifacts.

In this section, we shall present some other interpolations methods, that can be viewed as local approximations of the Fourier model. Our main concerns will be the locality of the interpolation (we shall impose that only a few values are used to define the image in the continuous neighborhood of a point) and, of course, the accuracy of the interpolate compared to the sinc interpolation.

2.2.1   Characterization

As we saw at the beginning of this chapter, a linear translation-invariant interpolation can be written

$$v(x) = \sum_{k \in \mathbb{Z}} u(k)\, \varphi(x - k). \tag{2.15}$$

Here we suppose that u is defined on $\mathbb{Z}$. All that follows could be presented in a periodic framework, as we did for the discrete Fourier model.

As we mentioned above, for computational reasons it is interesting to impose that the interpolating function (above, φ) has a small support (say included in $[-10, 10]$). This yields a distinction between direct and indirect local interpolations :

- a **direct local interpolation** is given by (2.15), where φ has a small support;

- an **indirect local interpolation** writes

$$v(x) = \sum_{k \in \mathbb{Z}} c(k)\, \psi(x - k), \tag{2.16}$$

where ψ has a small support. Since the interpolation is supposed to be exact — that is, $v(k) = u(k)$ for any $k \in \mathbb{Z}$ —, we must have

$$u = c \star \psi_d, \tag{2.17}$$

where ψ_d is the discrete signal associated to ψ (i.e. $\psi_d(k) = \psi(k)$ for any $k \in \mathbb{Z}$). If we consider discrete signals as sums of Dirac masses, taking the Fourier Transform of (2.17) yields

$$\hat{u}(\xi) = \hat{c}(\xi) \cdot \hat{\psi}_d(\xi),$$

so that one must have

$$\forall \xi \in \mathbb{R}, \qquad \hat{\psi}_d(\xi) := \sum_{k \in \mathbb{Z}} \psi_d(k) e^{-ik\xi} \neq 0.$$

In that case, the function φ in (2.15) will be defined by

$$\varphi(x) = \sum_{k \in \mathbb{Z}} (\psi_d)^{-1}(k) \psi(x - k), \tag{2.18}$$

with

$$\widehat{(\psi_d)^{-1}}(\xi) = (\hat{\psi}_d(\xi))^{-1}.$$

Hence, an initial pre-processing (the computation of $c = (\psi_d)^{-1} \star u$) has to be realized before the interpolation formula (2.16) can be applied. Note that in general φ will not have, like ψ, a small support.

Let us now define the order of interpolation.

Definition 4 *The interpolation method is said to be of order[1] $n \in \mathbb{N}$ if for any signal $U \in L^2(\mathbb{R})$ of class C^∞, one has*

$$\|U - U_h\|_2 = \underset{h \to 0}{o}(h^n),$$

where U_h is the signal interpolated from the discrete samples $U(kh), k \in \mathbb{Z}$.

The interpolation order gives an idea of the accuracy of the method : the larger the order, the more accurate the interpolate.

We shall also define the **support** w, which is the length of the smaller interval containing $\mathrm{supp}(\varphi)$. It can be shown that $w \geqslant n + 1$, where n is the order of interpolation.

Last, one can be interested in the **regularity** of φ (φ can be C^p, C^0, or even discontinuous), since this will be the regularity of v as well. Notice that regularity is concerned as long as we have to compute derivatives of the interpolate : since derivatives with order larger than 2 or 3 are barely used for images, aiming at a high order of regularity is not a real issue.

---

[1]Some authors define the *approximation* order of the method and use the convention that the order is $n + 1$ in this case, by reference to the order of a numerical scheme. However, this is sometimes confusing, and we prefer to define an *interpolation* order in agreement with the definition of the order of a Taylor expansion. In particular, as we shall see later, our definition ensures that the order of a spline is equal to its interpolation order.

2.2.2 Direct interpolations

• The simplest interpolation is called the **nearest neighbor interpolation**. It is a direct interpolation associated to

$$\varphi(t) = \beta^0(t) = \left\{ \begin{array}{ll} 1 & \text{if} \quad -\frac{1}{2} \leqslant t < \frac{1}{2}, \\ 0 & \text{else.} \end{array} \right.$$

It has order 0, support 1, and no regularity (β^0 is discontinuous). Its effect is very simple : the continuous signal is obtained as a piecewise constant signal defined from the discrete samples. In two dimensions, it yields images that are constant on adjacent squares.

• A slightly more complex method is given by the **(bi)linear interpolation**, associated to

$$\varphi(t) = \beta^1(t) = \beta^0 \star \beta^0(t) = \left\{ \begin{array}{ll} 1 - |t| & \text{if} \quad -1 \leqslant t \leqslant 1, \\ 0 & \text{else.} \end{array} \right.$$

It has order 1, support 2, and regularity C^0. The interpolated signal is piecewise linear.

In two dimensions, the bilinear interpolation is defined in the square $[k, k+1] \times [l, l+1]$ from the four values $u(k, l)$, $u(k+1, l)$, $u(k, l+1)$ and $u(k+1, l+1)$. Indeed, one has

$$\forall (x, y) \in [k, k+1] \times [l, l+1], \tag{2.19}$$

$$\begin{aligned} \tilde{u}(x, y) &= (l+1-y)\Big((k+1-x)u(k, l) + (x-k)u(k+1, l)\Big) \\ &+ (y-l)\Big((k+1-x)u(k, l+1) + (x-k)u(k+1, l+1)\Big). \end{aligned}$$

Like the nearest neighbor interpolation, the bilinear interpolation is monotone : if u and u' are two discrete signals such that $u(k) \geqslant u'(k)$ for any $k \in \mathbb{Z}$, then the corresponding interpolates satisfy $v(x) \geqslant v'(x)$ for any $x \in \mathbb{R}$. This property is a consequence of the fact that $\varphi \geqslant 0$ for both methods.

• If we look for a more accurate interpolation, then we may consider the so-called **bicubic** interpolation, associated to **Keys' function** (see Figure 2.1)

$$\varphi_a(t) = \left\{ \begin{array}{ll} (a+2)|t|^3 - (a+3)t^2 + 1 & \text{if} \quad 0 \leqslant |t| < 1, \\ a|t|^3 - 5at^2 + 8a|t| - 4a & \text{if} \quad 1 \leqslant |t| < 2, \\ 0 & \text{if} \quad 2 \leqslant |t|. \end{array} \right.$$

This method depends on a parameter $a \in [-1, 0]$. Its order is 1 in general, and 2 when $a = -1/2$. Its regularity is C^1, and its support is $w = 4$. It is very popular, probably because it is very simple to implement (this is a direct interpolation). However, we shall see that the cubic spline, an indirect interpolation function which has the same support ($w = 4$), is more accurate (order 3) and more regular (C^2). Thus, cubic spline interpolation should be preferred to Keys' bicubic interpolation in general.

2.2.3 Spline interpolation

The generalization of the nearest neighbor (β^0) and bilinear (β^1) interpolation can be realized with the theory of B-splines.

Definition 5 *The* **B-spline** *of order n,* β^n*, is defined recursively by*

$$\beta^0 = 1_{[-\frac{1}{2}, \frac{1}{2}[} \qquad \text{and} \qquad \beta^n = \beta^0 \star \beta^{n-1}.$$

Figure 2.1: Keys function for $a = -1/2$.

An important point to notice is that for $n \geqslant 2$, the interpolation associated to β^n is **indirect**, since we do not have any more $\beta^n(k) = 0$ for all $k \in \mathbb{Z}^*$. It is no longer monotone either. The regularity of β^n is C^{n-1}, and its support is minimal ($w = n + 1$). The function β^n can also be defined by

$$\beta^n(t) = \sum_{k=0}^{n+1} \frac{(-1)^k (n+1)}{k!(n+1-k)!} \cdot \max\left(0, \frac{n+1}{2} + t - k\right)^n.$$

The Fourier Transform of β^0 is simply

$$\widehat{\beta^0}(\xi) = \int_{\mathbb{R}} \beta^0(x) e^{-ix\xi} \, dx = \int_{-1/2}^{1/2} e^{-ix\xi} \, dx = \frac{\sin(\xi/2)}{\xi/2}.$$

Since $\beta^n = \beta^0 \star \beta^{n-1}$, this implies that

$$\forall n \in \mathbb{N}, \qquad \widehat{\beta^n}(\xi) = \left(\frac{\sin(\xi/2)}{\xi/2}\right)^{n+1}. \tag{2.20}$$

Since we are considering an indirect interpolation method (associated to $\psi = \beta^n$), we have to prove that it is well defined, i.e. that the convolution by ψ_d is invertible. This is the purpose of

Lemma 1 *Let $b^n = \sum_{k \in \mathbb{Z}} \beta^n(k)\delta_k$. Then, one has*

$$\forall \xi \in \mathbb{R}, \qquad \widehat{b^n}(\xi) \neq 0.$$

Proof :
    Taking the Fourier Transform of $b^n = \beta^n \cdot \Pi_1$ yields $\widehat{b^n} = \widehat{\beta^n} \star \Pi_{2\pi}$, that is, thanks to (2.20),

$$\forall \xi \in \mathbb{R}, \qquad \widehat{b^n}(\xi) = \sum_{p \in \mathbb{Z}} \left(\frac{\sin(\pi p + \xi/2)}{\pi p + \xi/2}\right)^{n+1} = \sum_{p \in \mathbb{Z}} (-1)^{p(n+1)} \left(\frac{\sin(\xi/2)}{\pi p + \xi/2}\right)^{n+1}.$$

If n is odd, all the terms of the sum are nonnegative. If $\xi \in [0, 2\pi[$, the term corresponding to $p = 0$ is positive, so that by periodicity $\widehat{b^n}(\xi) > 0$ for all $\xi \in \mathbb{R}$. The case "n even" is left as an exercise to the reader (indication : group the term p and $-p$ in the sum above). $\qquad \square$

Corollary 1 *The direct interpolating function φ associated to $\psi = \beta^n$ is called the* **cardinal spline** *of order n, written β_{card}^n. It is defined by*

$$\beta_{card}^n(x) = \sum_{k \in \mathbb{Z}} (b^n)^{-1}(k) \beta^n(x - k), \tag{2.21}$$

or equivalently by its Fourier Transform,

$$\widehat{\beta_{card}^n}(\xi) = \frac{\widehat{\beta^n}(\xi)}{\widehat{b^n}(\xi)} = \frac{\left(\frac{\sin(\xi/2)}{\xi/2}\right)^{n+1}}{\sum_{p\in\mathbb{Z}}\left(\frac{\sin(\pi p+\xi/2)}{\pi p+\xi/2}\right)^{n+1}}. \tag{2.22}$$

Now, a natural question is : what happens when $n \to \infty$? One can prove easily that when conveniently rescaled, β^n tends to a Gaussian function (we shall see in a further chapter that this is a consequence of the Central Limit Theorem). More interesting is the fact that the **cardinal spline** of order n tends to the cardinal sine (in L^2). Hence, spline interpolation is, for infinite n, equivalent to Shannon reconstruction formula.

Theorem 4 *In the L^2 sense,*

$$\beta_{card}^n \quad \underset{n\to\infty}{\longrightarrow} \quad \text{sinc}.$$

Proof :

We assume in the following that n is odd (the case "n even" is left as an exercise to the reader). By Parseval's Theorem, we may as well prove the announced result by showing that

$$\widehat{\beta_{card}^n} \quad \underset{n\to\infty}{\longrightarrow} \quad 1_{[-\pi,\pi]}$$

in $L^2(\mathbb{R})$. To do this, we notice that for $\xi \notin \pi\mathbb{Z}$, the numerator of (2.22) does not vanish, so that we can write

$$\widehat{\beta_{card}^n}(\xi) = \frac{1}{1 + U^n(\xi)}$$

with

$$U^n(\xi) = \sum_{p=1}^{\infty}\left[\left(\frac{2\pi p}{\xi}+1\right)^{-n-1} + \left(\frac{2\pi p}{\xi}-1\right)^{-n-1}\right]. \tag{2.23}$$

- If $\xi \in]0,\pi[$, then $t := 2\pi/\xi > 2$ and

$$U^n(\xi) \leqslant 2\sum_{p=1}^{\infty}\left(\frac{2\pi p}{\xi}-p\right)^{-n-1} \leqslant 2\left(\frac{2\pi}{\xi}-1\right)^{-n-1}\sum_{p=1}^{\infty}p^{-n-1} \leqslant 4\,(t-1)^{-n-1}$$

using the fact that $2\sum_{p\geqslant 1}p^{-n-1} \leqslant \pi^2/3 \leqslant 4$. Hence,

$$\int_0^{\pi}\left(\frac{1}{1+U^n(\xi)}-1\right)^2 d\xi \leqslant \int_0^{\pi}(U^n(\xi))^2\,d\xi \leqslant \int_2^{+\infty}\left(4\,(t-1)^{-n-1}\cdot\frac{2\pi}{t^2}\,dt\right)^2 = \underset{n\to\infty}{O}\left(\frac{1}{n}\right).$$

- If $\xi \in]\pi,+\infty[\backslash\pi\mathbb{Z}$, then we deduce from (2.23) that

$$U^n(\xi) \geqslant \max_{p\geqslant 1}\frac{\xi^{n+1}}{(2\pi p-\xi)^{n+1}} \geqslant \left(\frac{\xi}{\pi}\right)^{n+1}.$$

Thus,

$$\int_{\pi}^{+\infty}\left(\frac{1}{1+U^n(\xi)}\right)^2 d\xi \leqslant \int_{\pi}^{+\infty}\left(\frac{1}{1+(\xi/\pi)^{n+1}}\right)^2 d\xi = \underset{n\to\infty}{O}\left(\frac{1}{n}\right).$$

From the 2 cases investigated above and the fact that $\widehat{\beta_{card}^n}$ is even, we deduce that

$$\int_{\mathbb{R}}\left(\widehat{\beta_{card}^n}-1_{[-\pi,\pi]}\right)^2 \underset{n\to\infty}{\longrightarrow} 0,$$

which proves the announced result. $\qquad\qquad\square$

2.3 Reduction

Signal or image reduction (zoom out) cannot be formalized as a problem of exact interpolation, since a loss of information is necessary. We can compare different approaches.

- A direct subsampling of the original signal is probably the worst method, since it will in general introduce severe aliasing.

- A hard frequency cutoff (Equation 1.20), that is a projection on Shannon condition, may be realized as a direct application of the Fourier model. This corresponds to the projection on band-limited signals with a given resolution. However, we saw in the previous chapter that such a hard frequency cutoff has an important drawback : since it corresponds to a sinc convolution in spatial domain, it produces ringing around edges.

- A more reasonable solution is to perform a smooth frequency cutoff, using a preliminary convolution of the signal before sampling. This convolution may be done with the Gauss function

$$G_\sigma(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma^2}}.$$

or with a *prolate* function, both kernels being optimal in a certain sense. The Gauss kernel optimizes the spatial/spectral concentration according to Heisenberg inequality :

Theorem 5 (Heisenberg inequality) *If $f \in L^2(\mathbb{R})$, then*

$$\int_{\mathbb{R}} t^2 f^2(t)\, dt \cdot \int_{\mathbb{R}} \xi^2 |\hat{f}(\xi)|^2\, d\xi \geqslant \frac{\pi}{2} \cdot \left(\int_{\mathbb{R}} f^2(t) \right)^2,$$

with equality if and only if $f(t) = \alpha e^{-\beta t^2}$ for some $\alpha \in \mathbb{R}$ and $\beta > 0$.

A prolate function is a function with a given (small) spatial support, that maximizes the concentration of its energy in a given spectral support. It yields in general better result than the Gauss kernel.

Such a convolution-based approach has a major drawback : it is not a projection, so that even a very smooth signal which could be directly subsampled (for example because it has been obtained by enlarging another signal) will be smoothed a little more.

- Last, as an approximation of the projection on Shannon condition, a L^2 projection into the space of B-splines with order n may be used. This method is probably the best : like the hard frequency cutoff, it is a projection (hence, it will not produce a systematic blur), but for small n it avoids ringing artifacts.

Definition 6 *The space of B-splines with order n and sampling step δ is*

$$S_\delta^n = \left\{ x \mapsto \sum_{k \in \mathbb{Z}} c_k\, \beta^n \left(\frac{x}{\delta} - k \right),\ (c_k) \in l^2(\mathbb{Z}) \right\}. \tag{2.24}$$

The orthogonal projection on S_δ^n can be evaluated using

Theorem 6 *The orthogonal (L^2) projection of a signal $u \in L^2(\mathbb{R})$ on S_δ^n is*

$$\tilde{u}(x) = \frac{1}{\delta} \sum_{k \in \mathbb{Z}} \prec u, \overset{\circ}{\beta^n} \left(\frac{\cdot}{\delta} - k \right) \succ \beta_{card}^n \left(\frac{x}{\delta} - k \right),$$

where the dual spline of order n is defined by

$$\overset{\circ}{\beta^n}(x) = \sum_{k \in \mathbb{Z}} (b^{2n+1})^{-1}(k)\beta^n(x-k). \tag{2.25}$$

Notice the similarity between (2.25) and (2.21).

Chapter 3

The morphological model

3.1 Morphological invariance

Human vision is quite insensitive to absolute light intensity : wearing sunglasses, for example, does not change our perception of the world and our ability to identify objects or people (see Figure 3.1). Indeed, this ability is mostly based on the analysis of relative contrast : we essentially notice the relative gradient of intensity between two regions, provided that they are adjacent. As remarked by Kanizsa [17] and other gestaltists, our perception of intensities is so much based on local contrast that in some situations, the effect of the local background can lead us to a wrong perception of the relative contrast between two distant regions (see Figure 3.2).

Figure 3.1: *Our ability to understand images is quite insensitive to their absolute distribution of intensity, provided that enough contrast is available : these two images, which have been taken from the same viewpoint but with different camera settings, would be described in the same way.*

This suggests that image analysis should be performed independently of any reasonable grey-level rescaling. This sound principle has been used extensively in Mathematical Morphology [20, 29, 21, 15], and is for that reason sometimes called *morphological invariance.*

3.2 Contrast change and histogram

Let us represent a grey-level[1] image by a function $u : \Omega \to \mathbb{R}$, where Ω is a subset of $\mathbb{R}^2$ and $u(\mathbf{x})$ measures the intensity (grey level) at point $\mathbf{x}$. In this part, we shall investigate the numerical point of view and Ω will be a finite set (typically, a rectangular grid), whose cardinality will be

---

[1]we shall not consider color images here, but what follows could be applied to the intensity channel of a color image in an appropriate color space [12].

Figure 3.2: *Our perception of intensities is based on local contrast : though the central square on the left is lighter than the one on the right, we cannot help seeing that it is darker because of the strong influence of the local background (experiment from Kanizsa [17]).*

written $|\Omega|$. A *contrast change* is a nondecreasing function $g : \mathbb{R} \to \mathbb{R}$, that transforms an image u into the image $g \circ u$, also written $g(u)$. We shall note $\mathcal{M}$ the set of all contrast changes, and $\mathcal{M}_0$ the set of increasing ones. If $g \in \mathcal{M}_0$, then changing u into $g(u)$ is an invertible operation, which means that we can define a relation of equivalence with

$$u \sim v \quad \Leftrightarrow \quad \exists g \in \mathcal{M}_0, \ v = g(u).$$

If we have $v = g(u)$ with $g \in \mathcal{M} \setminus \mathcal{M}_0$, then there may be a loss of information between u and v and we shall write $v \in C(u)$.

Definition 7 (repartition function) *Let $u \in \mathbb{R}^\Omega$. The repartition function of u is the non-decreasing function*

$$H_u(t) = \frac{|\{\mathbf{x} \in \Omega; u(\mathbf{x}) \leqslant t\}|}{|\Omega|},$$

naturally extended to $\bar{\mathbb{R}} = [-\infty, +\infty]$ by $H_u(-\infty) = 0$ and $H_u(+\infty) = 1$. If we define as well

$$H_u^-(t) = \frac{|\{\mathbf{x} \in \Omega; u(\mathbf{x}) < t\}|}{|\Omega|},$$

then the average repartition function of u is

$$\bar{H}_u(t) = \frac{1}{2}\big(H_u(t) + H_u^-(t)\big).$$

Definition 8 (histogram) *The histogram of an image $u \in \mathbb{R}^\Omega$ is the distributional derivative of its repartition function, that is,*

$$H_u' = \sum_{i=1}^{n} h_u(\lambda_i)\delta_{\lambda_i},$$

where $\lambda_1 < \lambda_2 < \ldots < \lambda_n$ is the set of values taken by u, δ_t the Dirac mass in t, and

$$h_u(t) = \frac{|\{\mathbf{x} \in \Omega, u(\mathbf{x}) = t\}|}{|\Omega|}.$$

For numerical grey-level images coded on 8 bits, the possible values of u belong to $\{0..255\}$, so that the histogram of u is often defined as the function $h_u : \{0..255\} \to \mathbb{N}$.

Notice that the images $H_u \circ u$ and $H_u^- \circ u$ measure the *rank* of u : if u is one-to-one (which means that $n = |\Omega|$), then

$$u(x) = \lambda_k \quad \Leftrightarrow \quad H_u \circ u(x) = \frac{k}{|\Omega|} \quad \Leftrightarrow \quad H_u^- \circ u(x) = \frac{k-1}{|\Omega|}.$$

3.3   Histogram change

Can we prescribe the histogram (or equivalently, the repartition function) of a given image by applying a contrast change ? If we change an image u into $v = g \circ u$ with $g \in \mathcal{M}_0$, the repartition function becomes $H_v = H_u \circ g^{-1}$. Since H_u is a piecewise constant function, it is not possible in general to ensure that we would have $H_v = L$ for some g, even if L satisfies the properties of a repartition function (that is, $|\Omega|L$ is an integer-valued nondecreasing function such that $\lim_{-\infty} L = 0$ and $\lim_{+\infty} L = 1$). However, the following theorem tells us how we should choose g in order to have $H_v \simeq L$, and that this optimal choice is the same for various measures of the "error" between H_v and L.

Theorem 7 *Let $F : \mathbb{R}^+ \to \mathbb{R}^+$ be an increasing function such that $F(0) = 0$ (cost function) and $L : \bar{\mathbb{R}} \to [0,1]$ a nondecreasing function (the ideal repartition function), such that $L : [a,b] \to [0,1]$ is one-to-one and onto. For any $u \in \mathbb{R}^\Omega$, the unique $v \in C(u)$ that minimizes the cost integral*

$$\int_{\mathbb{R}} F\left(|H_v(t) - L(t)|\right) \, dt \tag{3.1}$$

is independent of F and given by $v = L^{-1} \circ \bar{H}_u \circ u$.

Proof :

Let us write $\lambda_1 < \lambda_2 < \ldots < \lambda_n$ the set of ordered values taken by u on Ω. If $v = g(u)$ with $g \in \mathcal{M}$, then $g_k = g(\lambda_k)$ is a nondecreasing sequence. In all the following, we use the convenient notations $\lambda_0 = g_0 = -\infty$ and $\lambda_{n+1} = g_{n+1} = +\infty$. Since

$$\forall k \in \{0..n\}, \ \forall t \in [g_k, g_{k+1}[, \quad H_v(t) = H_u(\lambda_k),$$

we have to minimize

$$E((g_k)) = \sum_{k=0}^{n} \int_{g_k}^{g_{k+1}} F\left(|H_u(\lambda_k) - L(t)|\right) \, dt \tag{3.2}$$

over all non-decreasing sequences $(g_k)_{1 \leqslant k \leqslant n}$ (g_0 and g_{n+1} being set as specified above).

1. We first minimize $E((g_k))$ over all (nondecreasing or not) sequences $(g_k) \in \mathbb{R}^n$. If (g_k) is such a sequence, we can define

$$h_k = \begin{cases} a & \text{if} \quad g_k < a, \\ g_k & \text{if} \quad a \leqslant g_k \leqslant b, \\ b & \text{if} \quad b \leqslant g_k. \end{cases}$$

Since F is increasing, we have $E((h_k)) \leqslant E((g_k))$, and the inequality is strict if the sequences (g_k) and (h_k) differ. Hence, the minimizers of E can be sought among the sequences $(g_k) \in [a,b]^n$. Since this set is compact, E has at least one global minimizer. Let (g_k) be such a minimizer. Since E is differentiable on $\mathbb{R}^n$, (g_k) is a critical point of E and we have, deriving (3.2),

$$\forall k \in \{1..n\}, \quad \frac{\partial E}{\partial g_k} = 0 = F\left(|H_u(\lambda_{k-1}) - L(g_k)|\right) - F\left(|H_u(\lambda_k) - L(g_k)|\right).$$

As F is injective, this yields

$$\forall k \in \{1..n\}, \quad H_u(\lambda_{k-1}) = H_u(\lambda_k) \quad \text{or} \quad H_u(\lambda_{k-1}) - L(g_k) = -H_u(\lambda_k) + L(g_k),$$

and the first equation is impossible. Thus, we necessarily have for all k

$$g_k = L^{-1}\left(\frac{H_u(\lambda_k) + H_u(\lambda_{k-1})}{2}\right), \tag{3.3}$$

which proves that E has a unique minimizer over $\mathbb{R}^n$. Notice that this minimizer, given by (3.3), is a non-decreasing sequence.

2. Now let us come back to our original task : we want to minimize E over the set S of all non-decreasing sequences. As the unique global minimizer of E over $\mathbb{R}^n$ is a nondecreasing sequence, it is as well the unique minimizer of E over S. Hence, there exists a unique $v \in C(u)$ that minimizes (3.1), and it is given by $v(\mathbf{x}) = g_{k(\mathbf{x})}$, where $k(\mathbf{x})$ is the unique index such that $u(\mathbf{x}) = \lambda_{k(\mathbf{x})}$. One easily checks that such a v belongs to $C(u)$ (any non-decreasing sequence can be extended into a non-decreasing function), and that $v = L^{-1} \circ \bar{H}_u \circ u$ is an equivalent definition of v. □

3.4 Histogram equalization

A classical way to analyze an image u independently of its contrast is the so-called histogram equalization. It consists in applying a special contrast change that flattens the histogram, or equivalently, forces the repartition function to be as linear as possible. This amounts to the particular case $L : t \in [0, 1] \mapsto t$ of Theorem 7. Hence, we know that histogram equalization is performed by setting

$$v = \bar{H}_u \circ u,$$

and that this choice minimizes in particular all L^p norms ($1 \leqslant p < \infty$) of $H_v - L$. It is not difficult to prove that this optimality property remains true for the L^∞ norm, except that $v = \bar{H}_u \circ u$ is not the unique minimizer of $\|H_v - L\|_\infty$ in general. We leave it as an exercise to the reader to determine the set of functions v such that $\|H_v - L\|_\infty$ is minimal.

An example of histogram equalization is shown on Figure 3.3. On Figure 3.4, we can check the effect of this transform on the repartition function, which becomes linear up to the unavoidable quantization effects. The influence on the histogram is much less clear, but by convolving it with a test function (typically a unit mass positive smooth function with little support), we could see that it is essentially constant in the sense of distributions.

Figure 3.3: *Left : original image. Right : the same image after histogram equalization.*

Histogram equalization is not an appropriate solution for contrast-invariant image analysis. The first reason is that it is completely global : the result of histogram equalization may be strongly influenced by a small dark or bright region of the image. This is not perceptually relevant, since we know that we observe scenes by quickly focusing our attention on different parts, which means that we are in some way insensitive to *local* contrast changes. We shall see later what the main local contrast-invariant features of a grey level image are. Another reason why histogram equalization is sometimes questionable is that it may strongly magnify the noise

Figure 3.4: *Repartition functions (up) and histograms (down) before (left column) and after (right) histogram equalization.*

amplitude in poorly contrasted regions, and thus produce undesirable noise textures. Since histogram equalization does not allow (at least in the formulation we presented) a control of the noise amplification, it seems dangerous to try to use it as the input of an image analysis device. Last, we can see on Figure 3.4 that in the case of quantized images, histogram equalization results in a loss of information, because some initially different grey levels may have to be transformed into the same grey tone in order to keep the initial (finite) set of possible grey values. This can hardly be avoided since the number of quantization levels for the grey tones is always smaller than $|\Omega|$.

3.5 How to estimate a contrast change ?

Applying a contrast change to an image allows to match approximately (in the sense of Theorem 7) its histogram with any preset one. Now we could ask the opposite question : when two images u and v are supposed to differ only by a contrast change, how can we find it ? This problem arises in particular in the context of image comparison, image fusion and mosaicking, and may be used for equalization purposes in video for example.

3.5.1 Comparison of two histograms

The simplest way to estimate a contrast change is to notice that if $v = g(u)$, then $H_v \circ g = H_u$, so that a contrast change g_0 between u and v could be estimated by composing a pseudo-inverse of H_v with H_u (notice that g_0 would be nondecreasing by construction). However, this direct method is not satisfactory since it does not guarantee that in the end $g_0 \circ u$ will be similar to v in some sense, especially if the initial assumption $v = g(u)$ was not exactly true, which is to be expected in real situations. Indeed, the bihistogram of u and v, defined in the discrete case by

$$h_{u,v}(i,j) = \frac{|\{\mathbf{x} \in \Omega, \ (u(\mathbf{x}), v(\mathbf{x})) = (i,j)\}|}{|\Omega|},$$

is more likely to be concentrated *around* a certain curve $j = g(i)$ (see Figure 3.5), so that the assumption $v = g(u)$ has to be weakened.

Figure 3.5: *Left : the bihistogram $h(i,j)$ of the two images of Figure 3.1. Right : the conditional bihistogram $h(i,j)/\sum_k h(i,k)$. The concentration of the mass in a "thick curve" shows that up to some unavoidable error, the two images are essentially related to each other by a contrast change.*

3.5.2   Least Squares Minimization

As we noticed, it seems more relevant to base the estimation of g on the simultaneous distribution (bihistogram) of u and v rather than on the marginal distributions (classical histograms) of u and v. Since we need to allow some noise, error, etc. in the correspondence between u and v, we shall rather minimize the L^2 square difference $\|v - g \circ u\|_2^2 = \sum_{\mathbf{x} \in \Omega} \left(v(\mathbf{x}) - g \circ u(\mathbf{x})\right)^2$ than impose exactly $v = g \circ u$.

Proposition 4 *Given two images $u : \Omega \to \mathbb{R}$ and $v : \Omega \to \mathbb{R}$, there exists a unique discrete function $g : u(\Omega) \to \mathbb{R}$ that minimizes $\|v - g \circ u\|_2$. If $\lambda_1, \lambda_2, \ldots, \lambda_N$ are the distinct values taken by u and $\Omega_k = \{\mathbf{x}; u(\mathbf{x}) = \lambda_k\}$, then one has*

$$g(\lambda_k) = \bar{v}_k := \frac{1}{|\Omega_k|} \sum_{\mathbf{x} \in \Omega_k} v(\mathbf{x}).$$

Proof :

Let us write $g_k = g(\lambda_k)$. We need to minimize

$$
\begin{aligned}
E((g_k)) &= \sum_{\mathbf{x} \in \Omega} (v(\mathbf{x}) - g \circ u(\mathbf{x}))^2 \\
&= \sum_{\mathbf{x} \in \Omega} v(\mathbf{x})^2 + \sum_{k=1}^{N} \sum_{\mathbf{x} \in \Omega_k} (g \circ u(\mathbf{x}))^2 - 2v(\mathbf{x}) g \circ u(\mathbf{x}) \\
&= \sum_{\mathbf{x} \in \Omega} v(\mathbf{x})^2 + \sum_{k=1}^{N} \left(|\Omega_k| g_k^2 - 2g_k \sum_{\mathbf{x} \in \Omega_k} v(\mathbf{x}) \right) \\
&= \sum_{\mathbf{x} \in \Omega} v(\mathbf{x})^2 + \sum_{k=1}^{N} |\Omega_k| \left(g_k^2 - 2g_k \bar{v}_k \right) \\
&= \sum_{\mathbf{x} \in \Omega} v(\mathbf{x})^2 - \sum_{k=1}^{N} |\Omega_k| \bar{v}_k^2 + \sum_{k=1}^{N} |\Omega_k| \cdot (g_k - \bar{v}_k)^2,
\end{aligned}
$$

and the unique minimum of $E((g_k))$ is attained for $g_k = \bar{v}_k$. $\qquad\square$

3.5.3 Monotone projection

The function g defined in Proposition 4 is very easy to compute but it is not a contrast change in the sense we defined previously, since we did not impose g to be nondecreasing. If we add this constraint to the minimization of $\|v - g \circ u\|_2$, we obtain the following

Definition 9 (monotone projection) *Let $u, v \in \mathbb{R}^\Omega$. We call monotone projection of u on v (written $M_v(u)$) the unique $u' \in C(u)$ that minimizes $\|v - u'\|_2$.*

Proposition 5 *Let $\lambda_1 < \lambda_2 < \ldots < \lambda_N$ be the ordered values taken by u, and Ω_k and $\bar{v}_k$ defined as in Proposition 4. Then $u(\mathbf{x}) = \lambda_k \Leftrightarrow M_v(u)(\mathbf{x}) = g_k$, where (g_k) is the unique nondecreasing minimizer of*

$$
\sum_{k=1}^{n} |\Omega_k| \cdot (g_k - \bar{v}_k)^2. \tag{3.4}
$$

The proof is an immediate consequence of Proposition 4. Notice that the sequence (g_k) is unique because it is the minimizer of a strictly convex function on the convex set of nondecreasing sequences.

Finding the unique nondecreasing (g_k) that minimizes (3.4) is a problem called *isotonic regression*, and the solution can be easily computed in linear time. Indeed, the sequence (g_k) can be obtained from $\bar{v}_k$ by successive merging of the "pool adjacent violators" (PAV) [3]. It can be shown that such a merging process converges to the optimal (g_k), no matter what order is used to combine the constant parts [25]. Hence we can use the following algorithm :

$$k = 1$$
$$\text{for } i = 1..N \text{ do}$$
$$\quad \sigma_k = |\Omega_i|\bar{v}_i$$
$$\quad w_k = |\Omega_i|$$
$$\quad n_k = 1$$
$$\quad \text{while } \left(k > 1 \text{ and } \frac{\sigma_{k-1}}{w_{k-1}} > \frac{\sigma_k}{w_k} \right) \text{ do}$$
$$\quad\quad \sigma_{k-1} = \sigma_{k-1} + \sigma_k$$
$$\quad\quad w_{k-1} = w_{k-1} + w_k$$
$$\quad\quad n_{k-1} = n_{k-1} + n_k$$
$$\quad\quad k = k - 1$$
$$\quad \text{end}$$
$$\quad k = k + 1$$
$$\text{end}$$
$$\text{return } (g_k), \text{ defined by } g_k = \frac{\sigma_i}{w_i} \Leftrightarrow \sum_{j=1}^{i-1} n_j < k \leqslant \sum_{j=1}^{i} n_j$$

On Figure 3.6 are shown three estimates of the contrast change involved between the two images of Figure 3.1. Apart from having a precise variational definition, the estimate of g yielded by the monotone projection is more robust to the differences that may arise between u and v, which may be caused by moving objects, occlusions, quantization, noise, saturation, etc. This can be explained by the fact that the monotone projection, contrary to the direct method presented in Section 3.5.1, is based on the joint distribution of u and v, so that no preliminary assumption on the correspondence between u and v has to be made.


Figure 3.6: *Estimated contrast changes between the two images of Figure 3.1. Left : direct match of the histograms. Middle : unconstrained L^2 minimization. Right : monotone projection. Notice the similarity with the joint distribution of the two images, represented on Figure 3.5.*

3.6   The level-set decomposition

Since the histogram of an image is precisely what is modified when a contrast change is applied, what is the remaining invariant part of images ? The answer lies in geometry : the equivalence classes of

$$u \sim v \quad \Leftrightarrow \quad \exists g \in \mathcal{M}_0, \ v = g(u)$$

can be described with the level-set decomposition [14]. In order to introduce geometry more easily, we shall now use a continuous model for images. In the following, an image will be a function u from $\mathbb{R}^2$ to $\mathbb{R}$. Numerically, this function will be obtained by interpolating the discrete image. We shall investigate later the theoretical and numerical issues related to this interpolation process.

Definition 10 (level sets) *The (upper) level sets of an image* $u : \mathbb{R}^2 \to \mathbb{R}$ *are the sets*

$$\chi_\lambda(u) = \{\mathbf{x} \in \mathbb{R}^2; \ u(\mathbf{x}) \geqslant \lambda\},$$

where λ *is any real number.*

For any $g \in \mathcal{M}_0$, the level sets of u and $g \circ u$ are the same, and more precisely one has

$$\chi_\lambda(u) = \chi_{g(\lambda)}(g \circ u).$$

Hence, the level sets do not depend on the absolute image contrast. They also contain all the image information, as shown below and on Figure 3.7.

Proposition 6 *Without any regularity assumption on the image* $u : \mathbb{R}^2 \to \mathbb{R}$, *one has*

$$\forall \mathbf{x} \in \mathbb{R}^2, \qquad u(\mathbf{x}) = \sup\{\lambda \in \mathbb{R}; \ \mathbf{x} \in \chi_\lambda(u)\}.$$

Proof :
   (immediate)

$$\sup\{\lambda \in \mathbb{R}; \ \mathbf{x} \in \chi_\lambda(u)\} = \sup\{\lambda \in \mathbb{R}; \ \lambda \leqslant u(\mathbf{x})\} = u(\mathbf{x}). \qquad \square$$

Level sets clearly follow a nesting property :

$$\forall \lambda, \mu, \qquad \lambda \geqslant \mu \quad \Rightarrow \quad \chi_\lambda(u) \subset \chi_\mu(u). \tag{3.5}$$

Let us suppose that u is upper semi-continuous, which means that all its upper level sets are closed, which allows to define their connected components. Equation 3.5 implies that the connected components of the level sets of u form a tree structure. This structure is inherited by the level lines defined below.

Definition 11 (level lines) *Let* $u : \mathbb{R}^2 \to \mathbb{R}$ *be an upper-semicontinuous image. The level lines of* u *are the topological boundaries of the connected components of its level sets.*

If we make a stronger regularity assumption, we can prove that these level lines are curves :

Proposition 7 *Let* L *be a non-empty bounded level line of a* C^1 *image* $u : \mathbb{R}^2 \to \mathbb{R}$. *Then,* $u_{/L}$ *is constant, and if* Du *does not vanish on* L, L *is a Jordan curve of class* C^1.

Figure 3.7: *Some upper level sets (in white) of the left image of Figure 3.1.*

Proof :

By definition, L is a connected component of an upper level-set of u, say $\chi_\lambda(u)$. Since this set is closed (u is C^1 and consequently upper-semicontinuous), we have $u(\mathbf{x}) \geqslant \lambda$ for any $\mathbf{x} \in L$. Now for any $\mathbf{x}$ such that $u(\mathbf{x}) > \lambda$ we know by continuity of u that the inequality $u(\mathbf{x}) > \lambda$ is still satisfied in a neighborhood of $\mathbf{x}$, which means exactly that $\mathbf{x}$ cannot belong to L. Hence, $u_{/L} = \lambda$.

Now assume that Du does not vanish on L. For any $\mathbf{x} \in L$, by implicit function Theorem we know that the equation $u(\mathbf{y}) = \lambda$ defines a C^1 graph in some neighborhood B of $\mathbf{x}$, and each point $\mathbf{y}$ of this graph belongs to $\partial\chi_\lambda(u)$. Since this graph is connected and contains $\mathbf{x}$, it is contained in L, which implies that it is equal to $L \cap B$. We deduce that, as a connected closed set of the plane that is locally a C^1 graph everywhere, L is a C^1 Jordan curve. $\qquad\square$

Notice that the requirement that Du does not vanish on L is not a strong assumption since we know by Sard's Theorem that if u is C^1, then for almost every λ one has $u(\mathbf{x}) = \lambda \Rightarrow Du(\mathbf{x}) \neq 0$.

3.7   Numerical computation of the level-lines tree

As we mentioned above, we need an interpolation model to build the tree of level lines of an image. The simplest interpolation method we can use is the so-called nearest-neighbor interpolation. If $u(i,j)$ (i,j integers) represents the numerical image, then its interpolate is simply defined by

$$\forall(x,y) \in \left[i - \frac{1}{2}, i + \frac{1}{2}\right[\times \left[j - \frac{1}{2}, j + \frac{1}{2}\right[, \qquad \tilde{u}(x,y) = u(i,j).$$

In this model, the continuous image is constant on each 1×1 square patch centered at the location of the known samples. The level lines of such an image are made of a concatenation of horizontal and vertical unit length segments, and depend on a convention of connectedness (4 or 8-connectedness).

The piecewise constant model does not yield very precise estimate of the level lines of a numerical image, and we can see on Figure 3.8 that much better results are obtained with the bilinear interpolation, defined by

$$\forall(x,y) \in [i, i+1] \times [j, j+1], \tag{3.6}$$

$$\tilde{u}(x,y) \;=\; (j+1-y)\Big((i+1-x)u(i,j)+(x-i)u(i+1,j)\Big)$$
$$+\;(y-j)\Big((i+1-x)u(i,j+1)+(x-i)u(i+1,j+1)\Big).$$

In each 1×1 square with integer coordinates of the image domain, $\tilde{u}$ is a bilinear function, and its level lines are hyperboles with horizontal and vertical axes. These pieces of hyperbole connect together at the boundaries of the 1×1 squares, but in general with a discontinuity of the tangent.

Figure 3.8: *Some level lines of the left image of Figure 3.1. Left column : nearest neighbor interpolation. Right column : bilinear interpolation. The finest scales (down) reveal the superiority of the bilinear interpolation over the zero order model. Another important improvement would probably be realized with higher order splines.*

The previous models correspond to spline interpolation of order 0 and 1 respectively. For the spline interpolation of order n (see [33] for example), we know that level lines along which Du does not vanish would be of class C^{n-1}. This suggests that we should use larger values of n, for example $n = 3$, which corresponds to cubic splines. Values above 3 are risky in that they may introduce undesirable ringing in the interpolated image, since the spline interpolation converges to sinc interpolation as n tends to infinity.

To compute at the same time all the level lines of an image (for a given set of thresholds λ), Pascal Monasse and Frederic Guichard have proposed a fast algorithm called the *Fast Level Set Transform*. We shall not describe this algorithm here, but the reader may refer to [24] for further details. The last version of this algorithm is available in the current release of the public software MegaWave2 (`http://www.cmla.ens-cachan.fr/~megawave`). It can deal with nearest-neighbor or bilinear interpolation.

The bilinear model has an evident computational interest : the level lines are computed explicitly as graphs in each 1×1 square and no new grey level can be created inside the square.

Like nearest-neighbor interpolation, bilinear interpolation is monotone, in the sense that all local strict extrema of the interpolated image belong to the set of initial samples. This property has some importance when all level lines of an image have to be computed at the same time. As far as we know, no algorithm exists at this time to compute the tree of the level lines of a numerical image in the context of cubic spline interpolation. This remains an interesting challenge, since such an algorithm could significantly improve the precision of the level-set decomposition of an image.

3.8 Scalar contrast-invariant local features

Let $u : \mathbb{R}^2 \to \mathbb{R}$ be C^∞. We would like to find local features of u that are invariant under contrast changes. We can do this by looking for differential operators T such that $T.u(\mathbf{x}) = T.g \circ u(\mathbf{x})$ for any smoothly increasing g. If the order of T is 1, then T is a function of Du, and since $D(g \circ u)(\mathbf{x}) = g' \circ u(\mathbf{x}) \cdot Du(\mathbf{x})$, we can see that Tu is necessarily a function of the orientation of u, defined when $Du(\mathbf{x}) \neq 0$ by

$$\text{orientation}(u)(\mathbf{x}) = \frac{Du(\mathbf{x})}{|Du(\mathbf{x})|}.$$

Moreover, one can check that orientation$(u)(\mathbf{x})$ is nothing but the normal to the level line of u that passes through $\mathbf{x}$.

Let us now consider the second order case. We shall write $(\mathbf{p}, \mathbf{q}) \in \mathbb{R}^2 \times \mathbb{R}^2 \mapsto D^2 u(\mathbf{p}, \mathbf{q})$ for the bilinear form associated to $D^2 u$, the 2×2 symmetric matrix representing the second order derivative of u. If T is a second order differential operator, then Tu depends only on Du and $D^2 u$. If we consider this representation in the local orthogonal basis $(Du, Du^\perp)$ given by the gradient (we assume that $Du(\mathbf{x}) \neq 0$), then we obtain three terms (all functions evaluated at point $\mathbf{x}$) :

- the *curvature* of u, defined by

$$\text{curv}(u) = \frac{1}{|Du|^3} D^2 u(Du^\perp, Du^\perp).$$

This corresponds to the curvature of the level line of u that passes through $\mathbf{x}$. If we note $\text{div} = \frac{\partial}{\partial x} + \frac{\partial}{\partial y}$, then an equivalent definition is

$$\text{curv}(u) = \text{div}\left(\frac{Du}{|Du|}\right).$$

- the *anticurvature* of u, defined by

$$\text{anticurv}(u) = \frac{1}{|Du|^3} D^2 u(Du, Du^\perp).$$

This is the curvature of the *gradient line* of u that passes through $\mathbf{x}$ (contrary to level lines, which are orthogonal the gradient everywhere, gradient lines are parallel to it).

- the *Canny operator*

$$\text{canny}(u) = \frac{1}{|Du|^3} D^2 u(Du, Du).$$

$|Du|\text{canny}(u)$ is nothing but the second derivative of u along gradient lines, that is to say in the direction where it varies the most quickly. Canny operator has been first studied in the context of edge detection [4], since it vanishes at the point where a contour profile reaches

an inflexion point (hence, solving $D^2u(Du, Du) = 0$ yields in principle a good localization of edges in images). It is also involved in the differential characterization of Absolutely Minimizing Lipschitz Extensions (AMLE) [2, 16], that can be used to interpolate strongly quantized images [6].

Thanks to the $1/|Du|^3$ normalization term, those three terms are invariant under the transform $u \mapsto \lambda u$. However, only the curvature and the anticurvature are invariant under contrast changes. This comes from the fact that

$$D^2(g \circ u)(\mathbf{p}, \mathbf{q}) = (g' \circ u) \cdot D^2u(\mathbf{p}, \mathbf{q}) + (g'' \circ u) \cdot \prec Du, \mathbf{p} \succ \prec Du, \mathbf{q} \succ .$$

Hence, every second-order contrast-invariant differential operator can be written as a function of the orientation, curv and anticurv operators.

Chapter 4

Local operators

4.1 Discretization of differential operators

How to translate a differential operator $u \mapsto F(u, Du, D^2u, \ldots)$ on a digital image $u : \mathbb{Z}^2 \to \mathbb{R}$? One possibility is to consider v, the Shannon interpolate of u, and to transpose F onto an operator $Tu = F(v, Dv, D^2v, \ldots)$. For example, the partial derivative $F(u, Du) = \frac{\partial u}{\partial x}$ would be translated that way into

$$
\begin{aligned}
Tu(i,j) &= \sum_{(k,l) \in \mathbb{Z}^2} u(k,l) \mathrm{sinc}'(i-k) \mathrm{sinc}(j-l) \\
&= \sum_{k \in \mathbb{Z}^*} u(k,j) \frac{(-1)^{i-k}}{i-k}.
\end{aligned}
$$

Such a definition is not satisfactory because it yields a strongly non local definition of a local operator. Apart from this major drawback, we also know that Shannon interpolation may suffer from ringing phenomena that will be magnified by the derivation process. At this stage, we have two possibilities : either we use a spline interpolation model instead of Shannon reconstruction, or we address the problem from the other side, by trying to understand local discrete operators as approximate differential operators. The latter solution is now considered.

We shall focus on translation-invariant operators, also called **filters**. Thus, a discrete filter is simply a map $T : \mathbb{R}^{\mathbb{Z}^2} \to \mathbb{R}^{\mathbb{Z}^2}$ such that

$$
\forall (\alpha, \beta) \in \mathbb{Z}^2, \qquad T \circ \tau_{\alpha,\beta} = \tau_{\alpha,\beta} \circ T,
$$

where $\tau_{\alpha,\beta}$ is the translation operator defined by

$$
\forall u \in \mathbb{R}^{\mathbb{Z}^2}, \ \forall (k,l) \in \mathbb{Z}^2, \qquad \left(\tau_{\alpha,\beta} u \right)(k,l) = u(k-\alpha, k-\beta).
$$

4.1.1 Consistency of a discrete operator

We first define the consistency of a discrete operator with a differential operator.

Definition 12 *Let $F(u, Du, \ldots, D^pu)$ be a differential operator of order $p \in \mathbb{N}$ ($F \neq 0$). We shall say that a discrete operator $T : \mathbb{R}^{\mathbb{Z}^2} \to \mathbb{R}^{\mathbb{Z}^2}$ is consistent with F if and only if there exists some $\alpha > 0$ such that*

$$
\forall u : \mathbb{R}^2 \to \mathbb{R}, \ u \in C^\infty, \ \forall (x,y) \in \mathbb{R}^2, \qquad \lim_{h > 0, h \to 0, kh \to x, lh \to y} h^{-\alpha} Tu_h(k,l) = F(u, Du, \ldots, D^pu)(x,y),
$$

where $u_h(k,l) = u(kh, lh)$.

Let us give some examples. The simplest local discrete operator consistent with $u \mapsto \frac{\partial u}{\partial x}$ is

$$Tu(k,l) = u(k+1,l) - u(k,l).$$

If more symmetry is required, one may use the centered finite difference scheme

$$T'u(k,l) = \frac{u(k+1,l) - u(k-1,l)}{2},$$

which is nothing but the convolution of T with the averaging horizontal kernel $\begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}$.

In many applications, the operator $F(u, Du) = |Du|$ needs to be translated numerically. There are several possibilities, even when the support is fixed. For a 2×2 support, one may use

$$Tu(k,l) = \sqrt{\frac{(b-a+d-c)^2 + (c-a+d-b)^2}{4}}, \tag{4.1}$$

with the convention that

$$a = u(k,l), \ b = u(k+1,l), \ c = u(k,l+1), \ d = u(k+1,l+1). \tag{4.2}$$

The estimation of $|Du|$ given by (4.1) satisfies

$$Tu(k,l) = |Du_b| \left(k + \frac{1}{2}, l + \frac{1}{2} \right),$$

where u_b is the bilinear interpolate of u. Note that Du_b is also related to the best L^2 local approximation of u by a linear form. However, the operator T defined in (4.1) does not characterize well constant images, since one may have $\forall k, l, \ Tu(k,l) = 0$ even for a non-constant image u (exercise : find all solutions). This cannot happen with

$$T'u(k,l) = \sqrt{\frac{(a-b)^2 + (d-c)^2 + (c-a)^2 + (d-b)^2}{2}}, \tag{4.3}$$

which is also consistent with $|Du|$. The latter estimate will be preferred in variational frameworks where $\int |Du|$ has to be minimized, since (4.1) will produce artifacts in general.

4.1.2 Consistency of a family of continuous operators

In a continuous framework, we can as well define the consistency of an operator $T : \mathbb{R}^{\mathbb{R}^2} \to \mathbb{R}^{\mathbb{R}^2}$, provided that a scaling parameter h has been introduced in the definition of T, yielding a family of operators $(T_h)_{h>0}$.

Definition 13 *The family of operators $(T_h)_{h>0}$ is consistent with F if and only if there exists some $\alpha > 0$ such that*

$$\forall u : \mathbb{R}^2 \to \mathbb{R}, \ u \in C^\infty, \ \forall (x,y) \in \mathbb{R}^2, \quad \lim_{h \to 0} h^{-\alpha} T_h(u)(x,y) = F(u, Du, \ldots, D^p u)(x,y).$$

Example. Let us investigate the asymptotic behavior of the continuous sup operator

$$S_h u(\mathbf{x}) = \sup_{|\mathbf{y}| \leqslant h} u(\mathbf{x} + \mathbf{y}).$$

If u is C^1, then one has

$$\forall \varepsilon > 0, \exists h > 0, |\mathbf{y}| \leqslant h \Rightarrow u(\mathbf{x}) + \prec Du(\mathbf{x}), \mathbf{y} \succ -\varepsilon h \leqslant u(\mathbf{x} + \mathbf{y}) \leqslant u(\mathbf{x}) + \prec Du(\mathbf{x}), \mathbf{y} \succ +\varepsilon h.$$

Taking the supremum for $|\mathbf{y}| \leqslant h$ on all sides, we obtain

$$u(\mathbf{x}) + h|Du(\mathbf{x})| - \varepsilon h \leqslant S_h u(\mathbf{x}) \leqslant u(\mathbf{x}) + h|Du(\mathbf{x})| + \varepsilon h,$$

that is

$$S_h u(\mathbf{x}) = u(\mathbf{x}) + h|Du(\mathbf{x})| + o(h).$$

Hence, the family $(S_h - id)_{h>0}$ is consistent with the differential operator $u \mapsto |Du|$.

4.1.3 Rank filters

Among discrete local operators, there is one class that has been extensively used in mathematical morphology : the rank filters.

Definition 14 *Let B be a finite subset of $\mathbb{Z}^2$. A discrete operator T_B is a rank filter if and only if there exists some function $f : \mathbb{R}^{|B|} \to \mathbb{R}$ such that*

$$\forall u \in \mathbb{R}^{\mathbb{Z}^2}, \forall \mathbf{x} \in \mathbb{Z}^2, \quad T_B u(\mathbf{x}) = f(\text{sort}(u(\mathbf{x} + \mathbf{y}), \mathbf{y} \in B)),$$

where

$$\text{sort}(t_{\sigma(1)}, t_{\sigma(2)}, \dots t_{\sigma(n)}) = (t_1, t_2, \dots t_n)$$

for any nondecreasing sequence $(t_i)_{i=1..n}$ and any permutation σ of $\{1..n\}$.

Intuitively, a rank filter is a discrete operator that does not refer explicitly in its definition to the position of the image samples. In a sense, it is a "non-geometric" operator, though we shall see later that some of them may be consistent with geometric operators.

Notice that contrast-invariant filters are not necessarily rank filters, as shown by the contrast-invariant operator

$$Tu(x, y) = u(x - \alpha, y).$$

Conversely, rank filters are not necessarily contrast invariant, as shown by the discrete averaging operator

$$M_B u(\mathbf{x}) = \frac{1}{|B|} \sum_{\mathbf{y} \in B} u(\mathbf{x} + \mathbf{y}). \tag{4.4}$$

Among examples of rank filters that are contrast-invariant, let us mention the inf and sup operators,

$$I_B u(\mathbf{x}) = \inf_{\mathbf{y} \in B} u(\mathbf{x} + \mathbf{y}) \tag{4.5}$$

$$S_B u(\mathbf{x}) = \sup_{\mathbf{y} \in B} u(\mathbf{x} + \mathbf{y}), \tag{4.6}$$

and the median operator that we shall define later in Section 4.3.

4.2 Isotropic positive linear filters

4.2.1 Asymptotics of an isotropic positive linear filter (continuous case)

By Riesz Theorem, a linear filter can be represented by a convolution, that is $Tu = \varphi \star u$. If $\varphi \geqslant 0$, the filter T is said to be positive. In the following, we shall mainly consider compactly supported filters, that is such that $\text{supp}(\varphi)$ is bounded. In this section, we shall see that any normalized isotropic positive compactly supported filter yields at second order the Laplacian operator.

Definition 15 *A linear filter $Tu = \varphi \star u$ is isotropic (up to order 2) if φ satisfies*

$$\int_{\mathbb{R}^2} x\varphi(x,y)\,dxdy = \int_{\mathbb{R}^2} y\varphi(x,y)\,dxdy = \int_{\mathbb{R}^2} xy\varphi(x,y)\,dxdy = 0 \tag{4.7}$$

$$\int_{\mathbb{R}^2} x^2\varphi(x,y)\,dxdy = \int_{\mathbb{R}^2} y^2\varphi(x,y)\,dxdy. \tag{4.8}$$

Moreover we shall say that T is normalized if

$$\int_{\mathbb{R}^2} \varphi(x,y)\,dxdy = 1, \tag{4.9}$$

Proposition 8 *Let $Tu = \varphi \star u$ be a normalized isotropic positive compactly supported linear filter, and consider the rescaled filter*

$$T_h u = \varphi_h \star u, \qquad \text{with} \qquad \varphi_h(\mathbf{x}) = h^{-1}\varphi(h^{-1/2}\mathbf{x}).$$

Then $(T_h - id)_{h>0}$ is consistent with $u \mapsto \sigma\Delta u$, where

$$\sigma = \frac{1}{4}\int_{\mathbb{R}^2} |\mathbf{x}|^2\varphi(\mathbf{x})\,d\mathbf{x}. \tag{4.10}$$

Proof :

First, we notice that φ_h satisfies (4.7), (4.8), (4.9), and that

$$\sigma_h := \frac{1}{4}\int_{\mathbb{R}^2} |\mathbf{x}|^2\varphi_h(\mathbf{x})\,d\mathbf{x} = \frac{1}{4}\int_{\mathbb{R}^2} h|\mathbf{y}|^2\varphi_h(h^{1/2}\mathbf{y})h\,d\mathbf{y} = h\sigma.$$

Moreover, if r is chosen such that $\text{supp}(\varphi) \subset B(0,r)$, then $\text{supp}(\varphi_h) \subset B(0, r\sqrt{h})$. Now, if u is C^2, then a Taylor expansion gives

$$\forall \varepsilon > 0, \exists h > 0, |\mathbf{y}| \leqslant r\sqrt{h} \Rightarrow -\varepsilon h \leqslant u(\mathbf{x}-\mathbf{y}) - u(\mathbf{x}) + \prec Du(\mathbf{x}), \mathbf{y} \succ -\frac{1}{2}[D^2 u(\mathbf{x})](\mathbf{y},\mathbf{y}) \leqslant \varepsilon h. \tag{4.11}$$

From (4.7), (4.8), and (4.9) we know that

$$T_h\left((x,y) \mapsto a + bx + cy + dx^2 + exy + fy^2\right) = a + 2(d+f)\sigma_h.$$

Hence, mutiplying each term of (4.11) by $\varphi_h(\mathbf{y})$ and integrating on $\mathbf{y} \in B(0, r\sqrt{h})$ yields

$$\forall \varepsilon > 0, \exists h > 0, \qquad -\varepsilon h \leqslant T_h u(\mathbf{x}) - u(\mathbf{x}) - \sigma h\Delta u(\mathbf{x}) \leqslant \varepsilon h,$$

which can be rewritten into

$$T_h u(\mathbf{x}) = u(\mathbf{x}) + \sigma h\Delta u(\mathbf{x}) + o(h). \tag{4.12}$$

This means exactly that the family $(T_h - id)_{h>0}$ is consistent ($\alpha = 1$) with the differential operator $u \mapsto \sigma\Delta u$. $\qquad\square$

Examples. Among simple examples of normalized isotropic positive compactly supported linear filters is the averaging operator

$$M_r u(\mathbf{x}) = \frac{1}{\pi r^2}\int_{|\mathbf{y}|\leqslant r} u(\mathbf{x}+\mathbf{y})\,d\mathbf{y}. \tag{4.13}$$

Let us take $T = M_r$ and investigate the asymptotics of $T_h = M_{r\sqrt{h}}$. The operator T satisfies (4.7), (4.8), (4.9), so that we only have to compute

$$\sigma = \frac{1}{4\pi r^2} \int_{|\mathbf{y}| \leqslant r} |\mathbf{x}|^2 \, d\mathbf{x} = \frac{1}{2r^2} \int_0^r \rho^2 \, \rho d\rho = \frac{r^2}{8}.$$

Thus, from Proposition (8) we deduce that the family $(T_h - id)_{h>0}$ is consistent with $u \mapsto \frac{r^2}{8} \Delta u$.

Another example, that we shall use later, is given by the Gauss kernel $\varphi = G_t$, where

$$G_t(\mathbf{x}) = \frac{1}{4\pi t} e^{-|\mathbf{x}|^2/4t}. \tag{4.14}$$

It is easy to show that G_t satisfies (4.7), (4.8), (4.9), and that the associated σ is equal to t. Even if this kernel is not compactly supported, the proof given above for Proposition 8 can be adapted with appropriate assumptions on u (for example u periodic, or $D^3 u$ bounded).

4.2.2 Asymptotics of an isotropic positive linear filter (discrete case)

Proposition 8 has a discrete counterpart, which is obtained by considering a distribution φ supported by $\mathbb{Z}^2$. Hence we can translate the definition and proposition above in the case of a discrete linear filter $Tu = A \star u$, where $A : \mathbb{Z}^2 \to \mathbb{R}$ is a finitely supported isotropic positive kernel.

Definition 16 *A discrete linear filter $Tu = A \star u$ is isotropic (up to order 2) if A satisfies*

$$\sum_{(k,l)\in\mathbb{Z}^2} kA(k,l) = \sum_{(k,l)\in\mathbb{Z}^2} lA(k,l) = \sum_{(k,l)\in\mathbb{Z}^2} klA(k,l) = 0 \tag{4.15}$$

$$\sum_{(k,l)\in\mathbb{Z}^2} k^2 A(k,l) = \sum_{(k,l)\in\mathbb{Z}^2} l^2 A(k,l) \tag{4.16}$$

Moreover we shall say that T is normalized if

$$\sum_{(k,l)\in\mathbb{Z}^2} A(k,l) = 1, \tag{4.17}$$

Proposition 9 *Let $Tu = A \star u$ be a discrete normalized isotropic positive finitely supported linear filter. Then, $T - id$ is consistent with $u \mapsto \sigma \Delta u$, where*

$$\sigma = \frac{1}{4} \sum_{(k,l)\in\mathbb{Z}^2} (k^2 + l^2) A(k,l). \tag{4.18}$$

Now we may study the consistency of the discrete averaging operator M_B defined in (4.4). If B is the discrete disc with radius r, that is $B = \{(k,l) \in \mathbb{Z}^2, \ k^2 + l^2 \leqslant r^2\}$, then $M_B - id$ is consistent with $\sigma \Delta u$, where

$$\sigma = \frac{1}{2|B|} \sum_{(k,l)\in B} k^2.$$

For example, if $B = \{(-1,0), (1,0), (0,0), (0,1), (0,-1)\}$ then $M_B - id$ is consistent with $\frac{1}{5} \Delta u$.

4.2.3   Deblurring with Gabor filter (sharpen)

Any isotropic linear kernel is, at first order, equivalent to a Gauss kernel, since they both have the same asymptotic expansion when conveniently rescaled. This yields an interesting application to image deblurring. As we shall see in Chapter 6, image deblurring can be formulated as follows : given an observed blurry and noisy image u_0, recover the physical image u that generated u_0, knowing that

$$u_0 = k \star u + n,$$

where k is the blur kernel (Point Spread Function) and n a noise term.

Suppose that the noise term is negligible (which is not always true !). If k is unknown, but positive and isotropic, it can be put under the form $k = C\varphi_h$, where φ is isotropic, normalized, and $\sigma(\varphi) = 1$. Now if h is reasonably small, we can write, according to (4.12),

$$\varphi_h \star u = u + h\Delta u + o(h),$$

which yields

$$u = \varphi_h \star u - h\Delta u + o(h) = \frac{1}{C}u_0 - h\Delta u + o(h),$$

and

$$\Delta u_0 = C\Delta(\varphi_h \star u) = C\varphi_h \star \Delta u = C\Delta u + o(1),$$

so that in the end,

$$Cu = u_0 - h\Delta u_0 + o(h).$$

Now, applying (4.12) to the Gauss kernel G_t, we get

$$G_t \star u_0 = u_0 + t\Delta u_0 + o(t),$$

so that we can approximate the ideal image Cu with

$$u_0 - \frac{h}{t}(G_t \star u_0 - u_0).$$

This is the Gabor (or sharpen) filter,

$$u_0 \mapsto u_0 + \lambda(u_0 - G_t \star u_0). \tag{4.19}$$

It has a simple interpretation : since $G_t \star u_0$ is a blurred version of u_0, the image $u_0 - G_t \star u_0$ contains the "details" of u_0, so that we increase the sharpness of u_0 by adding them (with a weight λ) to u_0. An example of deblurring with Gabor filter is shown on Figure 4.1

4.3   The median filter

4.3.1   Definition

Let us now focus on one of the most popular filter of mathematical morphology, the median filter. Its definition is the same in the continuous and in the discrete case, but the hypotheses to be made are different. In the following, an image is a function $u : \Omega \to \mathbb{R}$ and a structuring element is a bounded set $B \subset \Omega$. In the continuous case, $\Omega = \mathbb{R}^2$, B is compact, and u is assumed to be measurable. In the discrete case we simply have $\Omega = \mathbb{Z}^2$. In both situations, the measure associated to Ω (Lebesgue measure or counting measure) is written $|\cdot|$.

Figure 4.1: *Deblurring with Gabor filter. Left : original image. Right : processed image.*

Definition 17 *Let u be an image and B a structuring element. For $\mathbf{x} \in \Omega$ and $\lambda \in \mathbb{R}$, we define*

$$I_\lambda(\mathbf{x}) = \{\mathbf{y} \in B,\ u(\mathbf{x} + \mathbf{y}) < \lambda\} \qquad \text{and} \qquad S_\lambda(\mathbf{x}) = \{\mathbf{y} \in B,\ u(\mathbf{x} + \mathbf{y}) > \lambda\}.$$

Then, the median of u with respect to B is the (set) function

$$\mathrm{med}_B\, u\ :\ \mathbf{x} \mapsto \left\{ \lambda \in \mathbb{R},\ |I_\lambda(\mathbf{x})| \leqslant \frac{|B|}{2} \text{ and } |S_\lambda(\mathbf{x})| \leqslant \frac{|B|}{2} \right\}.$$

Since $\lambda \mapsto |I_\lambda(\mathbf{x})|$ is nondecreasing and $\lambda \mapsto |S_\lambda(\mathbf{x})|$ is nonincreasing, the set $\mathrm{med}_B\, u(\mathbf{x})$ is a closed interval, that can be written

$$\mathrm{med}_B\, u(\mathbf{x}) = [\mathrm{med}_B^-\, u(\mathbf{x}), \mathrm{med}_B^+\, u(\mathbf{x})].$$

If the function u is continuous on $\Omega = \mathbb{R}^2$, then $\mathrm{med}_B^-u(\mathbf{x}) = \mathrm{med}_B^+\, u(\mathbf{x})$ and $\mathrm{med}_B\, u(\mathbf{x})$ is (improperly but conveniently) considered as a function from $\mathbb{R}^2$ to $\mathbb{R}$. The same situation occurs ($\mathrm{med}_B^-\, u = \mathrm{med}_B^+\, u$) when $\Omega = \mathbb{Z}^2$ and $|B|$ is odd. In the following we shall assume that either u is continuous or $|B|$ is odd to simplify, so that $\mathrm{med}_B u$ will be considered as a real-valued function.

4.3.2 Properties

Proposition 10 *The median is a monotone contrast-invariant operator, that is*

$$\forall u, v, \qquad u \leqslant v \quad \Rightarrow \quad \mathrm{med}_B\, u \leqslant \mathrm{med}_B\, v,$$

$$\text{and} \qquad \forall u,\ \forall g \in \mathcal{M}_0, \qquad \mathrm{med}_B\, g(u) = g(\mathrm{med}_B\, u).$$

Equivalently, the median filter can be viewed as a "stack filter", since it commutes with thresholdings. If we define the operator T_λ by

$$T_\lambda u(\mathbf{x}) = \begin{cases} 1 & \text{if} \quad u(\mathbf{x}) \geqslant \lambda, \\ 0 & \text{else}, \end{cases}$$

then we have

$$T_\lambda(\mathrm{med}_B u) = \mathrm{med}_B(T_\lambda u).$$

Being contrast-invariant, the median filter is also a rank filter, like the mean, the inf and the sup operators mentioned above.

4.3.3 Consistency

Now we raise the natural question of the consistency of the median filter when B is a disc.

Theorem 8 ([14]) *Let $u : \mathbb{R}^2 \to \mathbb{R}$ be C^3. Then, at any point $\mathbf{x}$ where $Du(\mathbf{x}) \neq 0$, one has*

$$\operatorname{med}_{B(0,r)} u(\mathbf{x}) = u(\mathbf{x}) + \frac{r^2}{6} |Du(\mathbf{x})| \operatorname{curv}(u)(\mathbf{x}) + \underset{r \to 0}{o}(r^2). \tag{4.20}$$

This theorem proves that the family $(\operatorname{med}_{B(0,rh)} - id)_{h>0}$ is consistent with

$$F(u, Du, D^2 u) = \frac{r^2}{6} |Du| \operatorname{curv}(u).$$

4.4 Inf, sup, and the Canny operator

In this section, we study the infinitesimal behavior of the inf and sup operators defined by

$$I_h u(\mathbf{x}) \quad = \quad \inf_{|\mathbf{y}| \leqslant h} u(\mathbf{x} + \mathbf{y}) \tag{4.21}$$

$$S_h u(\mathbf{x}) \quad = \quad \sup_{|\mathbf{y}| \leqslant h} u(\mathbf{x} + \mathbf{y}) \tag{4.22}$$

As we saw in Section 4.1.2, the family $(S_h - id)_{h>0}$ is consistent with $u \mapsto |Du|$. By changing u into $-u$, we can see as well that $(I_h - id)_{h>0}$ is consistent with $u \mapsto -|Du|$. Now we refine these estimates to understand the limit behavior of $\frac{I_h + S_h}{2} - id$.

Proposition 11 *Let us define*

$$m_h u(\mathbf{x}) = \frac{I_h u(\mathbf{x}) + S_h u(\mathbf{x})}{2}.$$

The family $(m_h - id)_{h>0}$ is consistent (for $Du \neq 0$) with

$$F(u, Du, D^2 u) = \frac{1}{2} |Du| \operatorname{canny}(u), \quad \text{where} \quad \operatorname{canny}(u) = \frac{1}{|Du|^3} D^2 u(Du, Du).$$

Proof :

We only need to prove that if u is C^2, and $Du(\mathbf{x}) \neq 0$, then

$$S_h u(\mathbf{x}) = u(\mathbf{x}) + h|Du(\mathbf{x})| + \frac{h^2}{2} [D^2 u(\mathbf{x})] \left(\frac{Du(\mathbf{x})}{|Du(\mathbf{x})|}, \frac{Du(\mathbf{x})}{|Du(\mathbf{x})|} \right) + o(h^2). \tag{4.23}$$

By changing u into $-u$, it will prove that

$$I_h u(\mathbf{x}) = u(\mathbf{x}) - h|Du(\mathbf{x})| + \frac{h^2}{2} [D^2 u(\mathbf{x})] \left(\frac{Du(\mathbf{x})}{|Du(\mathbf{x})|}, \frac{Du(\mathbf{x})}{|Du(\mathbf{x})|} \right) + o(h^2)$$

and the announced result will follow immediately. One has

$$\forall \varepsilon > 0, \exists h > 0, |\mathbf{y}| \leqslant 1 \Rightarrow u(\mathbf{x}) + A_h(\mathbf{y}) - \varepsilon h^2 \leqslant u(\mathbf{x} + h\mathbf{y}) \leqslant u(\mathbf{x}) + A_h(\mathbf{y}) + \varepsilon h^2,$$

where

$$A_h(\mathbf{y}) = h \prec Du(\mathbf{x}), \mathbf{y} \succ + \frac{h^2}{2} [D^2 u(\mathbf{x})](\mathbf{y}, \mathbf{y}).$$

Taking the supremum for $|\mathbf{y}| \leqslant 1$ on all sides yields

$$u(\mathbf{x}) + \sup_{|\mathbf{y}| \leqslant 1} A_h(\mathbf{y}) - \varepsilon h^2 \leqslant S_h u(\mathbf{x}) \leqslant u(\mathbf{x}) + \sup_{|\mathbf{y}| \leqslant 1} A_h(\mathbf{y}) + \varepsilon h^2,$$

that is

$$S_h u(\mathbf{x}) = u(\mathbf{x}) + \sup_{|\mathbf{y}| \leqslant 1} A_h(\mathbf{y}) + o(h^2).$$

On one hand, we have

$$\sup_{|\mathbf{y}| \leqslant 1} A_h(\mathbf{y}) \geqslant A_h(\mathbf{y}_0), \quad \text{where} \quad \mathbf{y}_0 = \frac{Du(x)}{|Du(x)|}. \tag{4.24}$$

On the other hand, for $|\mathbf{y}| \leqslant 1$ and some constant $C > 0$,

$$A_h(\mathbf{y}) \leqslant A_h(\mathbf{y}_0) + h \prec Du(\mathbf{x}), \mathbf{y} - \mathbf{y}_0 \succ + Ch^2 |\mathbf{y} - \mathbf{y}_0|,$$

since the function $\mathbf{y} \mapsto [D^2 u(\mathbf{x})](\mathbf{y}, \mathbf{y})$ is locally Lipschitz. Notice that $\prec Du(\mathbf{x}), \mathbf{y} - \mathbf{y}_0 \succ \leqslant 0$.
• If $\prec Du(\mathbf{x}), \mathbf{y} - \mathbf{y}_0 \succ \leqslant -\sqrt{h}$, then

$$A_h(\mathbf{y}) \leqslant A_h(\mathbf{y}_0) - h\sqrt{h} + O(h^2) \leqslant A_h(\mathbf{y}_0) + o(h^2).$$

• If on the contrary $\prec Du(\mathbf{x}), \mathbf{y} - \mathbf{y}_0 \succ > -\sqrt{h}$, then

$$\begin{aligned}
|\mathbf{y} - \mathbf{y}_0|^2 &= \prec \mathbf{y} + \mathbf{y}_0, \mathbf{y} - \mathbf{y}_0 \succ - 2 \prec \mathbf{y}_0, \mathbf{y} - \mathbf{y}_0 \succ \\
&= |\mathbf{y}|^2 - |\mathbf{y}_0|^2 - 2 \prec \mathbf{y}_0, \mathbf{y} - \mathbf{y}_0 \succ \\
&\leqslant -2 \prec \mathbf{y}_0, \mathbf{y} - \mathbf{y}_0 \succ \\
&< 2\sqrt{h}/|Du(\mathbf{x})|,
\end{aligned}$$

so that for $C' = C(2/|Du(\mathbf{x})|)^{1/2}$,

$$A_h(\mathbf{y}) < A_h(\mathbf{y}_0) + C'h^2 h^{1/4} \leqslant A_h(\mathbf{y}_0) + o(h^2).$$

In both cases, we have $A_h(\mathbf{y}) \leqslant A_h(\mathbf{y}_0) + o(h^2)$, which proves with (4.24) that

$$A_h(\mathbf{y}) = A_h(\mathbf{y}_0) + o(h^2),$$

from which (4.23) and Proposition 11 follows. $\qquad \square$

Exercise : study the consistency of $8M_h - 2m_h - 6id$, where M_h is the continuous averaging operator defined in (4.13).

Solution. We know that $M_h - id$ is consistent with $u \mapsto \frac{1}{8}\Delta u$. Since

$$\Delta u = \text{trace}(D^2 u) = |Du|(\text{canny}(u) + \text{curv}(u)),$$

we deduce that $8M_h - 2m_h - 6id$ is consistent (for $Du \neq 0$) with

$$F(u, Du, D^2 u) = \Delta u - |Du| \, \text{canny}(u) = |Du| \, \text{curv}(u).$$

Chapter 5

The iterative model

The differential operators (gradient, orientation, curvature, etc.) we mentioned in the previous chapters are not so easy to estimate, and finite differences schemes yield poor results in general : as noticed by Marr [22], edges would be found everywhere when characterized by local differential operators in numerical images. As any physical data, images present some random fluctuations (noise) that may corrupt the true signal, but this explanation is not satisfactory since the problem still persists for high quality images with a large signal-to-noise ratio. The main reason why local estimates cannot be used directly is probably to be sought in the multiscale nature of images. Indeed, most images contain details with various sizes, which cannot all be handled by a single local operator. Thus, a finite differences scheme will be essentially driven by the little amount of highest resolution details, and will poorly capture the information it is supposed to.

If images follow a multiscale structure, then one could think that image analysis should too. Behind this idea is the notion of scale-space, introduced by Witkin [36]. In this framework, a raw image $u_0(\mathbf{x})$ is to be seen as a collection of images $u(\mathbf{x}, t)$, where t is a non-negative scale factor. For $t = 0$, we impose that $u(\mathbf{x}, 0) = u_0(\mathbf{x})$, while the image $\mathbf{x} \mapsto u(\mathbf{x}, t)$ is simplified as the scale increases. This smoothing process allows to compute derivatives, as in the case of distributions for example, in the sense that it selects an appropriate scale of details that will drive the estimate of the derivatives. Then, differential local features like gradient, orientation, curvature, etc. are now depending on the location variable, $\mathbf{x}$, but also on the scale variable t.

5.1   The heat equation

5.1.1   Iterated isotropic positive linear filters

In Proposition 8, we saw a consistency result for isotropic positive filters. We now consider the iteration of such filters, and show that the asymptotic behavior of this process leads to the heat equation.

Theorem 9 (Iterated convolutions) *Let φ be a normalized isotropic positive kernel, and consider the rescaled kernel*

$$\varphi_h(\mathbf{x}) = h^{-1}\varphi(h^{-1/2}\mathbf{x}).$$

Then, as $n \to \infty$, $h \to 0$, while keeping $nh\sigma = t$, one has

$$\varphi_h^{\star n} := \varphi_h \star \varphi_h \star \ldots \star \varphi_h \longrightarrow G_t,$$

where G_t is the Gauss kernel defined in (4.14) and σ is defined in (4.10).

This theorem is a direct consequence of the Central Limit Theorem below. As we shall see in the next section, it implies that the iteration of any positive isotropic linear kernel leads, up

to a rescaling, to the heat equation

$$\frac{\partial u}{\partial t} = \Delta u. \tag{5.1}$$

This is to be related to the consistency of $(T_h - id)_{h>0}$ with $u \mapsto \sigma \Delta u$, where $T_h u = \varphi_h \star u$.

Theorem 10 (Central Limit) *Let (X_i) be a sequence of i.i.d (independent and identically distributed) random variables such that $E[X_i] = 0$ and $E[X_i^2] = \sigma^2$. Then, as $n \to \infty$, the law of the random variable*

$$Y_n = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} X_i$$

tends to the normal law $\mathcal{N}(0, \sigma)$, given by

$$P[Y \geqslant t] = \frac{1}{\sigma\sqrt{2\pi}} \int_t^{+\infty} e^{-u^2/2\sigma^2} \, du.$$

The generalization to random vectors is immediate : if (X_i) is a sequence of i.i.d random vectors of $\mathbb{R}^d$ with mean $E[X_i] = 0$ and covariance matrix $E[X_i^T X_i] = C$ (with C invertible), then as $n \to \infty$ the density of Y_n tends to

$$\mathbf{x} \mapsto \frac{1}{\sqrt{\det C \cdot (\sqrt{2\pi})^d}} e^{-\frac{1}{2}\mathbf{x}^T C^{-1} \mathbf{x}}.$$

Proof of Theorem 9 :

Let X_i be i.i.d random vectors of $\mathbb{R}^2$ with density φ. Since

$$E[X_i] = 0 \qquad \text{and} \qquad C = E[X_i^T X_i] = 2\sigma \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

we know by Central Limit Theorem that

$$Y_n = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} X_i$$

converges in law to the Normal Distribution with density

$$\frac{1}{2\pi\sqrt{\det C}} e^{-\frac{1}{2}\mathbf{x}^T C^{-1} \mathbf{x}} = G_\sigma(\mathbf{x}).$$

Let $\delta = \sigma/t$. Since $nh = \delta^{-1}$ we can rewrite

$$Y_n = \sqrt{\delta} \sum_{i=1}^{n} \sqrt{h} X_i,$$

which proves that the density of Y_n is $\mathbf{x} \mapsto \delta^{-1}\psi_n(\delta^{-1/2}\mathbf{x})$, with $\psi_n(\mathbf{x}) = \varphi_h^{\star n}(\mathbf{x})$. Thus, we obtain as expected

$$\varphi_h^{\star n}(\mathbf{x}) \xrightarrow[n\to\infty]{} \delta G_\sigma(\delta^{1/2}\mathbf{x}) = G_t(\mathbf{x}). \qquad \square$$

5.1.2   Solutions of the heat equation

In what follows, C denotes a rectangular domain of $\mathbb{R}^2$, and the lattice associated to C defines the notion of C-periodic functions. Let us consider the following problem : given $u_0 \in L^1(C)$, C-periodic, find $u(\mathbf{x}, t) : \mathbb{R}^2 \times \mathbb{R}^+ \to \mathbb{R}$, C-periodic with respect to $\mathbf{x}$, such that

- for all $t > 0$, u is C^2 and bounded on $\mathbb{R}^2 \times [t, +\infty[$,

- $\dfrac{\partial u}{\partial t} = \Delta u$ on $\mathbb{R}^2 \times]0, +\infty[$

- $\int_C |u(\mathbf{x}, t) - u_0(\mathbf{x})| \, d\mathbf{x} \to 0$ when $t \to 0$.

This linear evolution problem is **well-posed** in Hadamard's sense :

1. we have **existence** of a solution u for any initial datum u_0;

2. we have **uniqueness** of such a solution for any initial datum;

3. we have **stability** with respect to initial conditions : the application that associates to u_0 the solution at time t (i.e. $u(\cdot, t)$) is Lipschitz for all t.

With the assumptions we made on the domain (rectangular) and on the boundary conditions (periodicity), the solution is given explicitly by

$$u(\mathbf{x}, t) = (G_t \star u_0)(\mathbf{x}). \tag{5.2}$$

which comes from the fact that $\frac{\partial G_t}{\partial t} = \Delta G_t$ and that $G_t \to \delta_0$ as $t \to 0$. However, this formula is generally no longer valid if we replace the condition "u C-periodic" by a condition taken on the boundary of a domain $\Omega \subset \mathbb{R}^2$, typically

- a Dirichlet condition :
$$\forall \mathbf{x} \in \partial\Omega, \forall t, \quad u(\mathbf{x}, t) = u_0(\mathbf{x});$$

- or a Neumann condition :
$$\forall \mathbf{x} \in \partial\Omega, \forall t, \quad \frac{\partial u}{\partial \nu}(\mathbf{x}, t) = 0,$$

where $\frac{\partial u}{\partial \nu}$ means the derivative of u along the normal to the boundary of Ω.

This justifies the need for a **numerical scheme** to compute the solution in the general case.

5.1.3   An explicit numerical scheme

Numerically, we have the following discrete problem :

1. we measure u_0 in some points of a discrete grid;

2. we compute the evolution of these values according to a numerical algorithm;

3. we hope that the values obtained approximate well the corresponding measures that would be obtained for the theoretical solution u associated to u_0.

Let us consider the following scheme : given a discrete image $u_0(i, j)$, we define iteratively the sequence $u_n(i, j)$ by

$$u_{n+1}(i, j) = u_n(i, j) + s \Big(u_n(i+1, j) + u_n(i-1, j) + u_n(i, j+1) + u_n(i, j-1) - 4u_n(i, j) \Big), \tag{5.3}$$

where s is a real positive parameter. For $s = 1/5$, this amounts to the particular case of the averaging filter mentioned in Section 4.2.2.

The discrete function u_{n+1} is obtained from u_n by a convolution, that we shall write

$$u_{n+1} = u_n + A \star u_n, \qquad \text{with} \quad A = \begin{bmatrix} 0 & s & 0 \\ s & -4s & s \\ 0 & s & 0 \end{bmatrix}.$$

Be careful : this representation of A as an array is not a matrix representation, but the representation of a discrete 3×3 image defined on the domain $\{-1, 0, 1\}^2$.

5.1.4 Consistency

The first condition that the numerical scheme should satisfy is consistency. It indicates that the scheme is a discretization of the PDE (5.1). To check the consistency, we need to view $u_n(i, j)$ as the discretization of a continuous image flow. If we call h the corresponding space step and Δt the time step, we have

$$u_n(i, j) = u(ih, jh, n\Delta t). \tag{5.4}$$

Now we investigate the consistency by using Taylor expansions that will allow us to interpret the finite differences in (5.3) as partial derivatives. If u is C^3 we have

$$u_n(i + 1, j) = u(ih + h, jh, n\Delta t) = \left(u + h\frac{\partial u}{\partial x} + \frac{h^2}{2}\frac{\partial^2 u}{\partial x^2} \right)(ih, jh, n\Delta t) + O(h^3).$$

Performing the same operation with the other terms yields

$$(A \star u_n)(i, j) = sh^2 \Delta u(ih, jh, n\Delta t) + O(h^3).$$

Notice that we proved here exactly that the discrete operator A is consistent with the Laplacian operator, $u \mapsto \Delta u$. Since we also have

$$u_{n+1}(i, j) - u_n(i, j) = \Delta t \frac{\partial u}{\partial t}(ih, jh, n\Delta t) + O(\Delta t^2),$$

we can see that if we set

$$\Delta t = sh^2, \tag{5.5}$$

then (5.3) implies

$$\frac{\partial u}{\partial t} = \Delta u + O(\Delta t),$$

which means that the scheme is consistent with the PDE (5.1), at order 1 in time since the remaining term is $O(\Delta t^1)$.

5.1.5 Stability

The second condition we may require from a numerical scheme is stability. An evolution scheme is stable if for any initial datum (here, $u_0(i, j)$), the numerical solution $(u_n(i, j))$ remains bounded on any finite time interval, independently of the discretization steps. In the case of evolution equations defined by an elliptic operator (here, the Laplacian), we have a

- **sufficient condition** : the comparison principle (for a scheme that preserves constants).

   Definition 18 *An operator T satisfies the* **comparison principle** *if $u \leqslant v \Rightarrow Tu \leqslant Tv$ for all u and v.*

This principle guarantees the stability, since writing $a \leqslant u_0 \leqslant b$ (where a and b are two real numbers) yields $Ta \leqslant Tu_0 \leqslant Tb$, that is $a \leqslant u_1 \leqslant b$ since T preserves constants. We conclude as well for u_n by recurrence.

In the case of a linear operator, the comparison principle can be formulated more simply : it is equivalent to say that $u \geqslant 0$ implies $Tu \geqslant 0$. Note that this property is satisfied by the heat equation : if a function is nonnegative, it remains so when it evolves according to (5.1). As concerns the scheme (5.3), we have

$$u_{n+1} = \varphi \star u_n \qquad \text{with} \qquad \varphi = \begin{bmatrix} 0 & s & 0 \\ s & 1-4s & s \\ 0 & s & 0 \end{bmatrix},$$

so that this scheme satisfies the comparison principle if and only if $0 \leqslant s \leqslant 1/4$.

For linear schemes with constant coefficients and no boundary condition, we also have a

- **necessary condition** : Von Neumann criterion. It requires that the amplification of any pure wave by an iteration of the scheme has a modulus at most 1. This is clearly a Fourier condition : if we rewrite $u_{n+1} = \varphi \star u_n$ in the Fourier domain, we obtain $\hat{u}_{n+1} = \hat{\varphi} \cdot \hat{u}_n$, and Von Neumann criterion writes $|\hat{\varphi}(\xi)| \leqslant 1$ for all ξ. In the case of (5.3), we have

$$\hat{\varphi}(p,q) = 1 - 4s + s\left(e^{ip} + e^{-ip} + e^{iq} + e^{-iq}\right) = 1 + 2s\left(\cos(p) + \cos(q) - 2\right),$$

and $1 - 8s \leqslant \hat{\varphi}(p,q) \leqslant 1$ (this inequality is optimal). Thus, we can see that the condition $0 \leqslant s \leqslant 1/4$ previously obtained is both necessary and sufficient.

What happens if Von Neumann criterion is not satisfied ? Since the solution after n iterations is

$$\hat{u}_n = (\hat{\varphi})^n \cdot \hat{u}_0,$$

we see that the coefficients of the frequencies ξ such that $|\hat{\varphi}(\xi)| > 1$ will blow up, and the dominant behavior of u_n will be driven by the eigenvectors $\mathbf{x} \mapsto e^{i \prec \mathbf{x}, \xi \succ}$ corresponding to the eigenvalues $\hat{\varphi}(\xi)$ with highest modulus.

5.1.6 Convergence

The convergence of a numerical scheme means that the function u defined from (5.4) converges to the solution of the continuous problem when the discretization steps (h and Δt) go to 0. This notion may be the most important of all, but it is generally not studied directly for linear problems, since it is a consequence of consistency and stability :

Theorem 11 (Lax-Richtmyer) *For a well-posed linear evolution problem, any consistent finite difference scheme is convergent if and only if it is stable.*

5.1.7 Other numerical schemes

The numerical scheme we presented is the most simple one. It has the important property that it is explicit, which means that u_{n+1} is defined by a function of u_n, typically

$$u_{n+1} = u_n + F(u_n),$$

where F is consistent with the differential operator involved in the PDE, in our example the Laplacian. There exists another kind of schemes, the implicit schemes, for which u_{n+1} is defined implicitly from u_n, typically

$$u_{n+1} = u_n + F(u_{n+1}).$$

Implicit scheme are more difficult to implement (because one has to solve an implicit equation, typically to invert a linear system), but they are generally more stable.

The implicit version of the previous scheme for the heat equation is

$$u_{n+1} - u_n = u_{n+1} \star \begin{bmatrix} 0 & s & 0 \\ s & -4s & s \\ 0 & s & 0 \end{bmatrix}.$$

More generally, it is clear that for $0 \leqslant \theta \leqslant 1$,

$$u_{n+1} - u_n = [\theta u_{n+1} + (1-\theta)u_n] \star \begin{bmatrix} 0 & s & 0 \\ s & -4s & s \\ 0 & s & 0 \end{bmatrix}$$

is also consistent with (5.1), and we have the following cases :

$$
\begin{array}{ll}
\theta = 0 : & \text{explicit scheme} \\
\theta = \frac{1}{2} : & \text{Crank-Nicolson scheme (order 2)} \\
\theta = 1 : & \text{implicit scheme}
\end{array}
$$

It can be shown that these schemes are always stable for $1/2 \leqslant \theta \leqslant 1$ (no condition on s is required). We leave it as an exercise to the reader to find the stability condition for $0 < \theta < 1/2$.

5.2   Curvature-driven motions

The linear scale-space representation given by the heat equation has a strong drawback : it is not contrast invariant. This means that two perceptually equivalent images u and $g \circ u$ will not be represented in the same way across scales. Hence, we had rather span a scale-space by iterating a contrast-invariant operator (like the median) than a convolution. This leads us to morphological scale-spaces, described by curvature-driven motions.

In Section 4.3.3, we saw that $(\text{med}_{B(0,h)})_{h>0}$ is consistent with $u \mapsto \frac{1}{6}|Du|\,\text{curv}(u)$. This proves that the iterated median is consistent with the Mean Curvature Motion

$$\frac{\partial u}{\partial t} = |Du|\,\text{curv}(u). \tag{5.6}$$

The theory of *viscosity solutions*[1], which we shall not present here, allows to define weak solutions (only continuous) of this equation, and to guarantee the existence and uniqueness of the solution of the Cauchy problem $u(\mathbf{x}, 0) = u_0(\mathbf{x})$. Please refer to [7] for a historical account, and [8] for a recent and concise presentation. In this part, we shall consider only regular solutions, for which (5.6) holds in the classical (C^2) sense.

Since $\text{curv}(u) = |Du|^{-3}[D^2u](Du^\perp, Du^\perp)$, we can rewrite (5.6) under the form

$$\frac{\partial u}{\partial t} = [D^2u](\xi, \xi), \qquad \text{with} \qquad \xi = \frac{Du^\perp}{|Du|}. \tag{5.7}$$

---

[1]the term "viscosity" comes from a relaxation technique that was used to define weak solutions of this kind of PDE's before a direct formulation was found. The original name has been kept, although no "viscous term" is now explicitly involved in the theory of viscosity solutions.

This means that (5.6) performs an **anisotropic diffusion** along the local direction of the level lines. To illustrate this, we could prove that a local averaging along the segment $[\mathbf{x} - r\xi, \mathbf{x} + r\xi]$ would yield a consistent iterative scheme for (5.6). This is to be compared to the isotropic diffusion performed by the heat equation, that can be approximated by iterated isotropic averaging. This explains why the mean curvature motion tends to produce sharper images compared to the heat equation : since no diffusion occurs across edges (that is, level lines), the image is smoothed without getting blurry (see Figure 5.1).


Figure 5.1: *Comparison of the heat equation (middle) and the mean curvature motion (right) computed at scale $t = 3$ for the same original image (left). The anisotropic diffusion caused by the mean curvature motion produces much less blur than the isotropic diffusion realized by the heat equation.*

5.2.1 Geometric interpretation

More generally, curvature-driven motions can be written under the form

$$\frac{\partial u}{\partial t} = |Du| \, F\Big(\mathrm{curv}(u)\Big), \tag{5.8}$$

where F is odd, continuous and nondecreasing. The theory of viscosity solutions remains valid for (5.8), which holds in the classical sense when u is C^2. Notice that if $u(\mathbf{x}, t)$ is solution of (5.8), then so is $g(u(\mathbf{x}, t))$ for any $g \in \mathcal{M}$. This explains why the evolution of u_0 performed by (5.8) can be, as a contrast-invariant process, translated into an evolution of the level lines of u_0.

Proposition 12 *Let I be an open interval of $\mathbb{R}$, and $u(\mathbf{x}, t) : \mathbb{R}^2 \times I \to \mathbb{R}$ a smooth (C^2) solution of (5.8) such that for some real λ, one has*

$$\forall t \in I, \ \forall \mathbf{x} \in \mathcal{C}(t) = \{\mathbf{x}, \ u(\mathbf{x}, t) = \lambda\}, \qquad Du(\mathbf{x}, t) \neq 0$$

Then, $\mathcal{C}(t)$ (parameterized by $\mathbf{C}(\cdot, t)$) evolves according to the pointwise equation

$$\frac{\partial \mathbf{C}}{\partial t} = F(\kappa) \, \mathbf{N}, \tag{5.9}$$

where $\mathbf{N} = -\frac{Du(\mathbf{C}, t)}{|Du(\mathbf{C}, t)|}$ is the normal vector to $\mathcal{C}$ and $\kappa = \mathrm{curv}(u)(\mathbf{C}, t)$ its curvature.

Proof :

We first note that $\mathcal{C}(t)$ is a disjoint union of C^2 Jordan curves, that are level-lines of $u(\cdot, t)$. Let us assume that $\mathcal{C}(t)$ is made of a single curve (the case of multiple components is similar).

We can parameterize $\mathcal{C}(t)$ by $(p,t) \mapsto \mathbf{C}(p,t)$ with $(p,t) \in S^1 \times I$. Now, by definition of $\mathcal{C}(t)$, we have

$$\forall (p,t) \in S^1 \times I, \qquad u\left(\mathbf{C}(p,t),t\right) = \lambda,$$

and taking the derivative with respect to t yields

$$\prec \frac{\partial \mathbf{C}}{\partial t}, Du \succ + \frac{\partial u}{\partial t} = 0.$$

Since a curve evolution is defined up to any tangential velocity (which has an influence on the parameterization of the curve, but not on its geometry), we can assume that $\frac{\partial \mathbf{C}}{\partial t}$ is supported by the normal $\mathbf{N}$, which is itself equal to $-Du/|Du|$ so that

$$\frac{\partial \mathbf{C}}{\partial t} = \frac{1}{|Du|} \frac{\partial u}{\partial t} \mathbf{N} = F(\mathrm{curv}(u))\mathbf{N} = F(\kappa)\mathbf{N}. \qquad \square$$

In particular, Proposition 12 shows that (5.6) is a geometric evolution : its effect is that all level lines of the initial image evolve independently according to the equation

$$\frac{\partial \mathbf{C}}{\partial t} = \kappa \, \mathbf{N}. \tag{5.10}$$

This is why (5.10) and, by extension, (5.6), is called the mean curvature motion. The term "mean curvature" comes from the fact that the generalization of (5.10) in higher dimensions is a normal evolution of a hypersurface at a speed proportional to its *mean* curvature (that is, the average of all principal curvatures).

5.2.2   Iteration of morphological filters

As proven by (4.20), the median operator is consistent with the mean curvature motion (5.6). In fact, Guichard and Morel [14] proved that the iterated continuous median actually converges to the semi-group given by (5.6). More precisely, if $u_0 : \mathbb{R}^2 \to \mathbb{R}$ is Lipschitz, then

$$\mathrm{med}_{B(0,r)}^n u(\mathbf{x}) \to u(\mathbf{x},t) \quad \text{as} \quad r \to 0, n \to \infty \text{ and } \frac{nr^2}{6} \to t,$$

and the limit $u(\mathbf{x},t)$ satisfies (5.6) in the viscosity sense.

However, iterating the median operator is not a good way to solve (5.6), because the median generally blocks after several iterations. This is a general property of local contrast-invariant operators : since they only rely on the level lines of a discrete function in a small neighborhood, they have many fixed points. Another way to understand this is to think in terms of curve evolution : a level line on a discrete grid is constrained to move by at least one pixel or to remain still. In the case of curvature-driven motions, this constraint forces small curvatures (corresponding to small motions) to be viewed as zero curvature. For the discrete median filter, one can see for example that a discrete disc with radius at least $\sqrt{5}$ is invariant under a 3×3 median (see Figure 5.2).

Figure 5.2: *The 3×3 median shrinks a discrete disc with radius 2, but has no effect on a disc with radius $\sqrt{5}$.*

Although the motion of the level lines induced by the median operator can be theoretically (that is, asymptotically) interpreted as the PDE evolution

$$\frac{\partial \mathbf{C}}{\partial t} = F(\kappa)\, \mathbf{N}$$

with $F(\kappa) = \kappa$, numerically one rather observes a similar motion with $F(\kappa) = 0$ for $\kappa < \kappa_0$ and $F(\kappa) = \kappa$ for $\kappa \geqslant \kappa_0$, since, as we mentioned previously, handling the smallest curvatures would require an arbitrary large neighborhood B.

5.2.3   The affine scale-space and its geometric implementation

The median filter led us to the mean curvature motion, which is a special case of curvature-driven motions (5.8). Now, among all these possible curvature-driven motions, formulated either in a scalar (5.8) or in a geometrical (5.9) way, there is a very interesting special case, given by $F(\kappa) = \kappa^{1/3}$ (with the convention that $(-|s|)^{1/3} = -|s|^{1/3}$).

Proposition 13 *The affine scale-space, defined by*

$$\frac{\partial u}{\partial t} = |Du|\, \mathrm{curv}(u)^{1/3}, \tag{5.11}$$

is affine invariant : if u is a solution of (5.11) then so is $(\mathbf{x}, t) \mapsto u(A\mathbf{x}, t)$ for any $A \in SL(\mathbb{R}^2)$.

Proof :
   Let us write $v(\mathbf{x}, t) = u(A\mathbf{x}, t)$. We shall assume that u is C^2 for simplicity, but the proof is essentially the same in the viscosity sense : the computation below has simply to be applied to smooth test functions. In the following, we shall denote partial derivatives with indices, and write $\mathbf{x} = (x, y)$. Since $v_t(\mathbf{x}, t) = u_t(A\mathbf{x}, t)$, we may forget the time variable and we just have to prove that for any $A \in SL(\mathbb{R}^2)$, the operator $Hu = |Du|\, \mathrm{curv}(u)^{1/3}$ satisfies

$$\forall \mathbf{x}, \quad Hv(\mathbf{x}) = Hu(A\mathbf{x}). \tag{5.12}$$

Since

$$\mathrm{curv}(u) = \frac{u_y^2 u_{xx} - 2u_x u_y u_{xy} + u_x^2 u_{yy}}{(u_x^2 + u_y^2)^{3/2}},$$

this operator can be written

$$Hu = \left(u_y^2 u_{xx} - 2u_x u_y u_{xy} + u_x^2 u_{yy} \right)^{1/3}.$$

Now, as H is rotation-invariant (because so are $u \mapsto |Du|$ and $u \mapsto \mathrm{curv}(u)$), we can rotate if necessary the axes in order to have $A = \begin{pmatrix} a & 0 \\ b & \frac{1}{a} \end{pmatrix}$. Then,

$$
\begin{aligned}
v_x &= a u_x \circ A, \\
v_y &= b u_x \circ A + \frac{1}{a} u_y \circ A, \\
v_{xx} &= a^2 u_{xx} \circ A, \\
v_{xy} &= a b u_{xx} \circ A + u_{xy}, \\
v_{yy} &= b^2 u_{xx} \circ A + 2\frac{b}{a} u_{xy} \circ A + \frac{1}{a^2} u_{yy} \circ A,
\end{aligned}
$$

and finally,

$$(Hv)^3 = (a^2v_y^2 - 2abv_xv_y + b^2v_x^2)u_{xx} \circ A + (-2v_xv_y + 2\frac{b}{a}v_x^2)u_{xy} \circ A + \frac{1}{a^2}v_x^2u_{yy} \circ A = (Hu \circ A)^3. \qquad \square$$

The affine invariance is an important property for shape analysis. Indeed, when a planar object located at infinity (in practice at a distance much larger than the focal length of the optical device) is observed from two different viewpoints, we saw in Chapter 1 that the two images are related by an affine transform. This property will then be preserved in their scale-space representations, which may be interesting for shape recognition tasks.

How to compute numerically the solutions of (5.11) ? The affine and the morphological invariance are difficult to handle properly in a scalar formulation, so that we should compute all level lines of u and make them evolve according to the geometric counterpart of (5.11), that is

$$\frac{\partial \mathbf{C}}{\partial t} = \kappa^{1/3} \mathbf{N}. \tag{5.13}$$

Hopefully, there exists a simple scheme [23] to do this. It is based on the following property.

Theorem 12 *Let $s \mapsto \mathbf{C}(s)$ be a closed convex curve of class C^2, parameterized with arclength, and $\sigma > 0$. To each point of $\mathbf{C}(s)$, we associate $\mathbf{C}_\sigma(s)$, defined as the middle point of the chord $(\mathbf{C}(s-\delta), \mathbf{C}(s+\delta))$, where $\delta > 0$ is chosen in order that the area of the region enclosed by this chord and the piece of curve $\mathbf{C}/_{[s-\delta,s+\delta]}$ be equal to σ (cf. Figure 5.3). Then, one has*

$$\mathbf{C}_\sigma(s) = \mathbf{C}(s) + \omega\sigma^{\frac{2}{3}} \cdot \kappa(s)^{\frac{1}{3}} \mathbf{N}(s) + \underset{\sigma \to 0}{o}(\sigma^{\frac{2}{3}}), \tag{5.14}$$

where $\mathbf{N}(s)$ is the interior normal of the curve, $\kappa(s)$ its curvature, and $\omega = \frac{1}{2}\left(\frac{3}{2}\right)^{\frac{2}{3}}$.

Figure 5.3: *The affine erosion of a convex curve is the envelop (or, equivalently, the set of the middle-points) of the σ-chords of the curve, that is the chords that enclose together with the curve a region with area σ.*

Proof :

We use the fact that if $s \mapsto \Gamma(s)$ is a Jordan curve, then its (signed) area is given by Green's formula,

$$\frac{1}{2} \oint \left[\Gamma(s), \Gamma'(s)\right] ds$$

(we write $[u, v]$ for the determinant of two vectors u and v). In the case we consider, we have

$$\sigma = \frac{1}{2}F(s, \delta(s, \sigma)), \quad \text{where}$$

$$F(s,t) = \int_{s-t}^{s+t} \left[\mathbf{C}(h), \mathbf{C}'(h)\right] \, dh \; + \; \left[\mathbf{C}(s+t), \mathbf{C}(s-t) - \mathbf{C}(s+t)\right]$$

Taking the derivative with respect to t yields

$$
\begin{aligned}
\frac{\partial F}{\partial t}(s,t) &= \left[\mathbf{C}(s+t), \mathbf{C}'(s+t)\right] + \left[\mathbf{C}(s-t), \mathbf{C}'(s-t)\right] \\
&\quad + \left[\mathbf{C}'(s+t), \mathbf{C}(s-t)\right] - \left[\mathbf{C}(s+t), \mathbf{C}'(s-t)\right] \\
&= \left[\mathbf{C}(s+t) - \mathbf{C}(s-t), \mathbf{C}'(s+t) - \mathbf{C}'(s-t)\right].
\end{aligned}
$$

A Taylor expansion in $t = 0$ then gives

$$\frac{\partial F}{\partial t}(s,t) = \left[2t\mathbf{C}'(s) + o(t), 2t\mathbf{C}''(s) + o(t)\right] = 4t^2\kappa(s) + o(t^2),$$

which can be integrated to obtain

$$2\sigma = \frac{4}{3}\delta^3\kappa(s) + o(\delta^3).$$

Thus, whenever $\kappa(s) \neq 0$ we have

$$\delta(s,\sigma) = \left(\frac{3\sigma}{2\kappa(s)}\right)^{\frac{1}{3}} + o(\sigma^{\frac{1}{3}}),$$

and finally

$$
\begin{aligned}
\mathbf{C}_\sigma(s) &= \frac{1}{2}\left[\mathbf{C}(s-\delta) + \mathbf{C}(s+\delta)\right] \\
&= \mathbf{C}(s) + \frac{\delta^2}{2}\mathbf{C}''(s) + o(\delta^2) \\
&= \mathbf{C}(s) + \frac{1}{2}\left(\frac{3}{2}\right)^{\frac{2}{3}}\sigma^{\frac{2}{3}} \cdot \kappa^{\frac{1}{3}}(s)\,\mathbf{N}(s) + o(\sigma^{\frac{2}{3}}).
\end{aligned}
$$

The case $\kappa(s) = 0$ is left as an exercise to the reader. $\qquad\square$

Theorem 12 allows to build a simple iterative scheme to solve (5.13) with a polygonal curve. It consists in the iteration of the following three-steps process :

- decomposing a curve into convex parts;

- replace each component by the sequence of the middle point of each σ-chord such that one endpoint is a vertex of the polygonal curve;

- recompose the curve from the new convex parts.

This scheme is able to smooth all the level lines of an image (that is, several thousands of curves) in a couple of seconds. An example is given on Figure 5.4.

Figure 5.4: *The level lines of Figure 3.8, after some affine smoothing (Equation 5.13). Overall computation time : 1.5 second.*

Chapter 6

The variational model

6.1 Classical operators as minimizers

In the variational framework, problems are characterized by a functional (generally measuring some kind of reconstruction error), and solutions are defined as minimizers of this functional. We already used this paradigm to define optimal histogram changes in Theorem 7 by minimizing (3.1). Indeed, variational characterizations are everywhere and they generally allow interesting interpretations. The averaging operator,

$$M_B \, u(\mathbf{x}) = \frac{1}{|B|} \sum_{\mathbf{y} \in B} u(\mathbf{x} + \mathbf{y}),$$

for example, has a variational interpretation : $M_B \, u(\mathbf{x})$ is the unique minimizer of

$$t \mapsto \sum_{\mathbf{y} \in B} \left(u(\mathbf{x} + \mathbf{y}) - t \right)^2.$$

This is still true in the continuous case ($u : \mathbb{R}^2 \to \mathbb{R}$ measurable and B compact subset of $\mathbb{R}^2$) :

$$M_B \, u(\mathbf{x}) = \frac{1}{|B|} \int_B u(\mathbf{x} + \mathbf{y}) d\mathbf{y},$$

is the unique minimizer of

$$t \mapsto \int_B \left(u(\mathbf{x} + \mathbf{y}) - t \right)^2 d\mathbf{y}.$$

The median operator can be defined in a similar way : $\mathrm{med}_B u(\mathbf{x})$ is exactly the set of minimizers[1] of

$$t \mapsto \sum_{\mathbf{y} \in B} \left| u(\mathbf{x} + \mathbf{y}) - t \right|$$

in the discrete case ($u : \mathbb{Z}^2 \to \mathbb{R}$ and B finite subset of $\mathbb{Z}^2$), and of

$$t \mapsto \int_B \left| u(\mathbf{x} + \mathbf{y}) - t \right| d\mathbf{y}$$

in the continuous case. Like the averaging operator, the median filter is a particular case of M-filters, for which $T_B \, u(\mathbf{x})$ minimizes

$$t \mapsto \sum_{\mathbf{y} \in B} \psi \left(u(\mathbf{x} + \mathbf{y}) - t \right)$$

---

[1]As we noticed previously, this set is made of a unique element as soon as $|B|$ is odd (discrete case) or u is continuous (continuous case).

for some given function ψ (see [18]).

Some of the PDE's we encountered in the previous chapter can also be interpreted in a variational framework. The heat equation (5.1), for example, is the Euler equation (gradient descent) associated to the minimization of

$$E(u) = \frac{1}{2} \int |Du|^2.$$

Let us quickly justify this. If u and v belong to $H^1(\mathbb{R}^2) = W^{1,2}(\mathbb{R}^2)$ and $\lambda \in \mathbb{R}$, we have

$$\frac{d}{d\lambda}\left(E(u + \lambda v)\right)_{\lambda=0} = \frac{d}{d\lambda}\left(\frac{1}{2}\int Du^2 + \lambda \int \prec Du, Dv \succ + \frac{1}{2}\lambda^2 Dv^2\right)_{\lambda=0} = \int \prec Du, Dv \succ,$$

Now if we assume for example that u is compactly supported (note that a periodicity condition would do as well), we can integrate by parts and obtain

$$DE(u).v = -\int v \cdot \mathrm{div}(Du) = -\int v \cdot \Delta u,$$

that is, $\nabla E(u) = -\Delta u$. Hence, (5.1) is exactly the gradient descent of E.

Some PDE's, however, cannot be directly interpreted as the Euler equation of a minimization problem. The mean curvature motion (5.6), for example, cannot be put in divergence form. However, it can be shown that its geometric counterpart, (5.10), is the Euler equation associated to the minimization of the Euclidean perimeter of a curve.

6.2   Image Restoration

6.2.1   The $Ku + n$ model

The variational model is often used in the context of image restoration. A typical aim is to try to remove the blur and the noise of an observed image $u_0 : \mathbb{R}^2 \to \mathbb{R}$. This can be modeled by the equation

$$u_0 = Ku + n, \tag{6.1}$$

where u is the "original" image to be recovered. The term $Ku = k \star u$ means the convolution of u with the kernel k, which models a shift-invariant blur (diffraction, out of focus, etc.). We shall assume that k is known in the following. The last term n represents all fluctuations (noise, error) that may cause u_0 to be different from Ku. They can be due to the measure of intensities (thermal noise), to the quantization of grey levels, to the sampling process, etc. We have chosen an additive model, but other choices are possible (e.g. multiplicative noise, as studied by Osher and Rudin).

This unavoidable noise term explains why even if K is invertible (that is, $\hat{k}$ does not vanish), which is false in general, the restoration of u by a brute force deconvolution ($u = K^{-1}u_0$) cannot be efficient. Indeed, we can write in Fourier domain

$$\widehat{K^{-1}u_0} = \hat{u} + \frac{\hat{n}}{\hat{k}},$$

and the second term may become very large at points where $\hat{k}$ is very small, which produces oscillations that may dominate by far the true part of the reconstructed image.

6.2.2 Regularization

A solution to this kind of ill-posed problem is the so called Tikhonov regularization [32]. It consists in choosing, among all functions u such that

$$\|u_0 - Ku\| \leqslant \sigma \tag{6.2}$$

(σ being the prior variance of the noise), the smoothest one. Typically, the restored image u may be defined as the unique minimum of

$$\int |Du|^2 \tag{6.3}$$

under the constraint (6.2). Using Lagrange multipliers, one can see that this is equivalent to minimizing

$$E(u) = \int |Du|^2 + \lambda \int (u_0 - Ku)^2 \tag{6.4}$$

for some λ. In that precise case, the solution can be written explicitly in Fourier domain : the restoration process boils down to a particular case of Wiener filtering, defined by

$$\forall \xi \in \mathbb{R}^2, \qquad \hat{u}(\xi) = \frac{\hat{k}^*(\xi)}{|\hat{k}(\xi)|^2 + \frac{|\xi|^2}{\lambda}} \cdot \hat{u}_0(\xi)$$

(z^* is the conjugate complex of z).

6.2.3 Total Variation

The regularity term (6.3) does not provide a good model for images, because it favors too much smooth images. In particular, no discontinuous image is allowed by (6.3), which is not very realistic if we recall that the image formation process is driven by occlusions. In 1992, Rudin, Osher and Fatemi [26] proposed to use the L^1 norm of the gradient (instead of the L^2 norm) for the regularity term. Then, (6.3) is replaced by

$$\int |Du|, \tag{6.5}$$

and the natural space for u is no longer the Sobolev space H^1 but the space BV of L^1 functions with bounded variation, that is for which $|Du|$ (in the sense of distributions) is a measure with finite total mass [13, 37]. The energy (6.4) now becomes

$$E(u) = \int |Du| + \lambda \int (u_0 - Ku)^2, \tag{6.6}$$

and since no explicit solution exists in that case, an iterative scheme is required. We shall investigate this point later.

6.2.4 Link with the Bayesian framework

The variational formulations (6.4), (6.6), and in fact all energies that can be written as the sum of a fidelity term and of a regularity term, have an interesting interpretation in the Bayesian framework. The Maximum A Posteriori (MAP) permits to formalize image restoration as a statistical problem, where the posterior probability $P(u|u_0)$ has to be maximized with respect to u : we want to find the best estimate of u knowing the observed image u_0. The Bayesian rule, written

$$P(u|u_0) = \frac{P(u_0|u)P(u)}{P(u_0)}, \tag{6.7}$$

reduces this maximization to the maximization of $P(u_0|u)P(u)$, since $P(u_0)$ is merely a normalization constant once u_0 is given. The first term, $P(u_0|u)$, is the likelihood for u_0 being generated by a given "ideal" image u, while $P(u)$ is the *prior* probability of u, that is its probability before u_0 was observed. Now, by taking the logarithm, we can see that we have to minimize

$$E(u) = -\log P(u) - \log P(u_0|u),$$

and as in (6.4) and (6.6), the energy is a combination of a regularity term $(-\log P(u))$ and of a fidelity term $(-\log P(u_0|u))$. Conversely, the Total Variation model (6.6) can be interpreted in the Bayesian framework as a Maximum A Posteriori associated to the prior

$$P(u) = \frac{1}{Z} \exp\left(-\int |Du|\right)$$

and to a Gaussian noise model

$$P(u_0|u) = \frac{1}{Z'} \exp\left(-\lambda \int (u_0 - Ku)^2\right)$$

(Z and Z' are normalization constants here).

6.2.5   Links with the PDE model

As we mentioned above, finding the minimizer of (6.6) requires an iterative scheme, typically a gradient descent.

Proposition 14 *The gradient descent of (6.6) is given by*

$$\frac{\partial u}{\partial t} = \mathrm{div}\left(\frac{Du}{|Du|}\right) + 2\lambda K^*(u_0 - Ku). \tag{6.8}$$

Proof :

We simply need to compute the gradient of E. We have

$$
\begin{aligned}
D\left(u \mapsto \int |Du|\right).v &= \frac{d}{d\lambda}\left(\int |Du + \lambda Dv|\right)_{\lambda=0} \\
&= \frac{d}{d\lambda}\left(\int \sqrt{(Du + \lambda Dv)^2}\right)_{\lambda=0} \\
&= \int \frac{\prec Du, Dv \succ}{|Du|} \\
&= -\int v \cdot \mathrm{div}\left(\frac{Du}{|Du|}\right),
\end{aligned}
$$

the last equality being obtained by integration by parts. Similarly, we have

$$
\begin{aligned}
D\left(u \mapsto \int (u_0 - Ku)^2\right).v &= \frac{d}{d\lambda}\left(\int (u_0 - Ku - \lambda Kv)^2\right)_{\lambda=0} \\
&= \int 2Kv \cdot (u_0 - Ku) \\
&= -2\int v \cdot K^*(u_0 - Ku),
\end{aligned}
$$

where $K^*u = k^* \star u$ and $k^*(\mathbf{x}) = k(-\mathbf{x})$. The last equality comes from the fact that for $f \in L^2(\mathbb{R}^2)$,

$$
\begin{aligned}
\int (Kv)(\mathbf{x}) \cdot f(\mathbf{x})\, d\mathbf{x} &= \int \left(\int k(\mathbf{x} - \mathbf{y})v(\mathbf{y})\, d\mathbf{y} \right) \cdot f(\mathbf{x})\, d\mathbf{x} \\
&= \int \int k^*(\mathbf{y} - \mathbf{x})v(\mathbf{y})f(\mathbf{x})\, d\mathbf{y}\, d\mathbf{x} \\
&= \int v(\mathbf{y}) \cdot (K^*f)(\mathbf{y})\, d\mathbf{y}.
\end{aligned}
$$

Thus, the gradient of E writes

$$
DE(u).v = -\int v \cdot \left[\operatorname{div}\left(\frac{Du}{|Du|} \right) + 2\lambda K^*(u_0 - Ku) \right],
$$

from which (6.8) follows. $\qquad\square$

If λ is small, which means that the regularity term dominates the fidelity terms, we can see that the beginning (t small) of the gradient descent (6.8) will be essentially driven by the unconstrained evolution

$$
\frac{\partial u}{\partial t} = \operatorname{div}\left(\frac{Du}{|Du|} \right), \tag{6.9}
$$

which is a degenerate case of Perona-Malik model. The comparison between (6.8) and the PDE model could stop here, since scale-spaces rather focus on the intermediate states of the PDE evolution (that is, $u(t, \mathbf{x})$ for *all* t), whereas only the *asymptotic* state ($t = +\infty$) of (6.8) is relevant in the variational framework.

In fact, the relevant parameter in the variational framework is not the time t involved in the gradient descent, but the coefficient λ involved in the definition of the energy, as it is illustrated on Figure 6.1. As we shall see now, the weight $\frac{1}{2\lambda}$ plays a similar role as the scale variable t in a scale-space, since it rules the amount of simplification allowed in the restoration process. Let us rewrite (6.6) with $t = \frac{1}{2\lambda}$ as

$$
E_t(u) = 2t \int |Du| + \int (u_0 - Ku)^2
$$

and call $u(t, \mathbf{x})$ the minimizer of E_t. From (6.8), we know that $u(t, \mathbf{x})$ satisfies the PDE

$$
t \cdot \operatorname{div}\left(\frac{Du}{|Du|} \right) + K^*(u_0 - Ku) = 0
$$

with initial condition $u(0, \mathbf{x}) = u_0(\mathbf{x})$. If we derive this equation with respect to t, we obtain

$$
Hu + t(Hu)_t = K^*Ku_t \qquad \text{with} \qquad Hu = \operatorname{div}\left(\frac{Du}{|Du|} \right).
$$

Hence, for small values of t (that is, λ large), we have

$$
K^*Ku_t = \operatorname{div}\left(\frac{Du}{|Du|} \right) + \underset{t\to 0}{O}(t).
$$

In particular, in the pure denoising case (no blur, that is, $Ku = u$), we have

$$
u_t = \operatorname{div}\left(\frac{Du}{|Du|} \right) + \underset{t\to 0}{O}(t)
$$

which is, at first order, the degenerate case (6.9) of the Perona-Malik model we mentioned previously. We could do the same for the H^1 variational model (6.4), and we would obtain in a similar way

$$u_t = \Delta u + \underset{t\to 0}{O}(t),$$

which means that the heat equation achieves, at first order in t, the optimal denoising as formulated by (6.4).

Figure 6.1: *Image restoration (here, denoising) using the TV model. The relevant parameter is the weight λ of the fidelity term of (6.6), $t = \frac{1}{2\lambda}$ playing a similar role as the scale variable in scale-space representations. From left to right, then top to bottom : $t = 0$ (original image), $t = 5$, $t = 20$ and $t = 100$.*

6.3 Numerical issues

The computational minimization of (6.6) can be done either by using the gradient descent (6.8) [27], or by solving directly the equilibrium equation

$$\mathrm{div}\left(\frac{Du}{|Du|}\right) + 2\lambda K^*(u_0 - Ku) = 0$$

using iterative methods [34]. In any case, the total variation has to be relaxed, in order to extend the validity of the linearization process involved at each iteration : since the linearization of $\mathbf{x} \mapsto |\mathbf{x}|$ is inefficient around 0, a direct use of the total variation would force the time steps to remain very small and would prevent an actual convergence of the algorithm. This is why the approximate norm

$$|Du|_\varepsilon = \sqrt{\varepsilon^2 + |Du|^2}$$

(for some ε) is generally used for numerical purposes.

Let us describe in detail the gradient descent method. In the following, u is a discrete image defined on some rectangular domain Ω of $\mathbb{Z}^2$. Among the most simple estimates of $|Du|_\varepsilon$ is

$$|Du(k,l)|_\varepsilon = \sqrt{\varepsilon^2 + \frac{1}{2} \sum_{(i,j) \in B} \big(u(k+i, l+j) - u(k,l)\big)^2},$$

the sum being extended to $(i,j) \in B = \{(0,1), (0,-1), (1,0), (-1,0)\}$ (4 nearest neighbors). A mirror symmetry is applied along the boundary of Ω. Rather than looking for a numerical scheme for (6.8), we write a discrete version of (6.6),

$$E(u) = \sum_{k,l} |Du(k,l)|_\varepsilon + \lambda(Ku - u_0)^2(k,l), \tag{6.10}$$

and compute its gradient with respect to each value $u(k,l)$. We have

$$\frac{\partial |Du(k,l)|_\varepsilon}{\partial u(k,l)} = - \sum_{(i,j) \in B} \frac{u(k+i, l+j) - u(k,l)}{2|Du(k,l)|_\varepsilon},$$

while for $(i,j) \in B$,

$$\frac{\partial |Du(k+i, l+j)|_\varepsilon}{\partial u(k,l)} = - \frac{u(k+i, l+j) - u(k,l)}{2|Du(k+i, l+j)|_\varepsilon},$$

so that

$$\frac{\partial}{\partial u(k,l)} \sum_{k',l'} |Du(k',l')|_\varepsilon = -\kappa_\varepsilon(u)(k,l)$$

with

$$\kappa_\varepsilon(u)(k,l) = \sum_{(i,j) \in B} \frac{u(k+i, l+j) - u(k,l)}{2} \left(\frac{1}{|Du(k,l)|_\varepsilon} + \frac{1}{|Du(k+i, l+j)|_\varepsilon} \right).$$

Notice that $\kappa_\varepsilon(u)$ is a discrete approximation of the curvature of u, $\mathrm{curv}(u) = \mathrm{div}\left(\frac{Du}{|Du|}\right)$. Now, the iterative gradient descent of (6.10) writes

$$u^{n+1}(k,l) = u^n(k,l) + \delta^n \cdot v^n(k,l) \qquad \text{with} \qquad v^n(k,l) = \kappa_\varepsilon(u)(k,l) + 2\lambda K^*(u_0 - Ku)(k,l),$$

where the step δ^n is computed at each iteration in order to ensure that $E(u^n)$ decreases. The optimal step (that is, the value of δ that minimizes $E(u^n + \delta^n v^n)$ can be computed by Newton's method for example, but in practice the algorithm may converge more quickly when a more simple strategy based on a progressive reduction of δ^n is applied. One can use the following decision rule : set $\delta^n := \delta^{n-1}$ and while $E(u^n + \delta^n v^n) \geqslant E(u^n)$ do $\delta^n := \alpha \cdot \delta^n$, where $\alpha \in]0, 1[$ has been fixed in advance. We use the same strategy for ε. Since the limit case $\varepsilon = 0$ we are interested in cannot be reached directly, we minimize (6.10) iteratively and decrease geometrically the value of ε after each minimization, using the steady state obtained to initialize the new minimization process.

Since convergence may be difficult to reach for functionals involving the total variation, it is often very useful to visualize the evolution of the energy as a function of the number of iterations performed.

Chapter 7

Detection a contrario: an introduction

7.1 Computer Vision and human vision

The aim of Computer Vision is to extract data related to the physical world from images. The final step may or may not be the complete reconstruction and interpretation of the 3D-world from a set of planar images, but in all cases the detection of basic geometric features (like segments, corners, contours, etc.) is often required as a necessary preliminary step.

Human vision works in a similar way : as proven by the Gestaltists (Wertheimer [35], Kanizsa [17]), our perception of images is, before any cognitive process may be involved, driven by a reasonnable small set of principles called gestalts (alignment, color constancy, closedness, convexity, parallelism, etc.). These gestalts work recursively at all scales, and generally cooperate. For example, a black square on a white background is percieved as such thanks to the constant color, to the local alignment (the sides are straight lines), to the parallelism of opposite sides, to the convexity of the whole shape, etc.

Now, a main issue for automatic detection tasks is the setting of thresholds : if the detection method depends on one or several parameters, a region containing all detectable events has to be defined in the parameter space. If the model only involves one parameter, there is only a threshold to set, but when the number of parameters grows to 2,3,4, etc. there is in general no easy way to define the detection region. For example, the Hough Transform, The Mumford and Shah segmentation, and the Maximum A Posteriori model all suffer from the problem of setting parameters.

7.2 The a contrario model

The a contrario model has been developped a few years ago in the case of alignments. So far, it has been applied to about a dozen of other detection tasks, some of them being gestalts (contrasted edges and boundaries, similarity of a scalar quality with uniform or decreasing distribution, alignment of isolated points, clusters, constant width, vanishing points, corners) and some of them being more generic (mode of a histogram, detection of point matches in a stereo pair). In the following, we introduce some bases of the a contrario model and discuss two main cases ; the detection of segments (alignment of points sharing the same orientation) and the detection of contrasted edges (closed or not). To present the a contrario model, we first refer to a principle due to Helmholtz.

7.2.1 Helmholtz Principle

One of the most striking fact about human vision is our ability to detect unlikely geometric coincidences. If we draw many non-overlapping circles and one square on a sheet of paper, we shall immediatly distinguish an isolated element (the square) on an homogeneous background (the circles). Following Helmholtz, this may even be stated as a basic principle for human vision.

Helmholtz Principle. *Human vision is based on the following mechanism : we assume a priori that object qualities are randomly and independently distributed, and we organize objects into groups when this assumption is not realized.*

7.2.2 Meaningful events

Let us consider some abstract objects O_1, O_2, ... O_M, and suppose that for each object O_i we can compute some measure $\mu(O_i)$. We want to set a threshold in order to detect high values of μ.

First, we can apply Helmholtz Principle and assume a random independent distribution on the parameters that define μ. In the following, we shall write $\tilde{O}_i$ for the random object associated to this distribution, and O_i its observed realization.

Then, we could choose the detection thresholds $\mu_0(i)$ in order to ensure that

$$\mathrm{P}\left[\exists i,\ \mu(\tilde{O}_i) \geqslant \mu_0(i)\right] \leqslant \varepsilon. \tag{7.1}$$

Here, ε controls the reliability of the test (the smaller ε, the more reliable the test). Such a definition ensures that a false alarm occurs by chance with a probability smaller than ε.

However, the left term of (7.1) is in general very difficult to estimate, because the $\mu(\tilde{O}_i)$ have no reason to be independent. This is why the following criterion will be preferred.

Definition 19 *The number of false alarms associated to the object O_i is*

$$NFA(O_i) = M \cdot \mathrm{P}\left[\mu(\tilde{O}_i) \geqslant \mu(O_i)\right]. \tag{7.2}$$

The object O_i is said to be ε-meaningful if $NFA(O_i) \leqslant \varepsilon$.

This criterion raises, for each value of i, an optimal threshold $\mu_0(i)$, which is the smallest value of μ that satisfies

$$\mathrm{P}\left[\mu(\tilde{O}_i) \geqslant \mu\right] \leqslant \frac{\varepsilon}{M}.$$

Such a definition ensures that the expected number of false alarms is smaller than ε. Indeed, if we set

$$\chi_i = \begin{cases} 1 & \text{if} \quad NFA(O_i) \leqslant \varepsilon, \\ 0 & \text{if} \quad NFA(O_i) > \varepsilon, \end{cases}$$

then the expected number of false alarms is

$$\mathrm{E}\left[\sum_i \tilde{\chi}_i\right] = \sum_i \mathrm{E}\left[\tilde{\chi}_i\right] = \sum_i \mathrm{P}\left[NFA(\tilde{O}_i) \leqslant \varepsilon\right] = \sum_i \mathrm{P}\left[\mu(\tilde{O}_i) \geqslant \mu_0(i)\right] \leqslant \sum_i \frac{\varepsilon}{M} = \varepsilon.$$

7.2.3   A toy example : birthday dates

The advantage of using expectations instead of probabilities can be illustrated by the following example. Suppose that among a group of N persons, you observe that two of them have the same birthday date (modulo the year). Is that a meaningful event, that is, could this have happened by chance or is it likely that a cause has to be found (presence of twins for example) ?

According to Helmholtz principle, we should assume that all birthday dates have the same probability $1/365$ (we neglect leap years). Now, we have two possibilities.

1. Following (7.1), we compute the probability that at least 2 persons have the same birthday date, which is

$$1 - 1 \cdot \left(1 - \frac{1}{365}\right) \cdot \left(1 - \frac{2}{365}\right) \cdots \left(1 - \frac{N-1}{365}\right). \tag{7.3}$$

   Then, in order to determine when this probability is less than ε, we need a tabulation. (e.g: $\varepsilon = 1/2$ gives $N \geqslant 22$).

2. Following rather (7.2), we write $O_1, O_2, \ldots O_M$ all couples (k, l) with $1 \leqslant k < l \leqslant N$ (hence $M = N(N-1)/2$), and set $\mu(O_i) = 1$ if the two birthday dates corresponding to O_i are equal, $\mu(O_i) = 0$ else. Now, the event $\mu(O_i) = 1$ is ε-meaningful if

$$NFA(O_i) := \frac{N(N-1)}{2} \cdot \frac{1}{365} \leqslant \varepsilon. \tag{7.4}$$

   This criterion is much simpler than (7.3), since it writes approximately $N \leqslant 27\sqrt{\varepsilon}$.

7.3   Detection of alignments

7.3.1   Model

We now come to a real gestalt : the detection of segments in an image. Let u be a $N \times N$ image. In each point $\mathbf{x}$ of the image plane, we compute the direction $\theta(\mathbf{x}) \in S^1 = \mathbb{R}/2\pi\mathbb{Z}$ of the level line passing through $\mathbf{x}$ (that is, the direction of $Du^\perp$). This requires the estimation of the gradient of u, which we realize with a 2×2 scheme. Notice that when the gradient is zero, θ is not defined. In the following, we keep the convention that any test performed on θ fails in this case.

Now we consider $O_1, O_2, \ldots O_M$ all discrete segments of the image plane, defined by two distinct points. Thus, we have $M = N^2(N^2 - 1) \simeq N^4$, and to each segment O_i we can associate a direction $\theta_i \in S^1$. Now, let l be the maximum number of independent points of O_i, that is, the maximum number points far enough from each others to ensure that no dependency occurs in the estimation of the orientations. Amonf these l points, we may count the number k of such points $\mathbf{x}$ for which

$$\text{dist}(\theta(\mathbf{x}), \theta_i) \leqslant p\pi, \tag{7.5}$$

the distance being taken on the circle S^1, and p being a fixed parameter (typically $p = 1/16$). This way, we measure how many points have their orientation aligned (up to the precision p) with the segment itself. Then, given k and l, we would like to decide when the segment is meaningful (in the sense that k is unexpectedly large).

To that aim, we can apply Helmholtz Principle to the orientation field, and assume that $\theta(\mathbf{x})$ are independent random variables, uniformly distributed in S^1. We then have

$$\mathrm{P}\left[\text{dist}(\tilde{\theta}(\mathbf{x}), \theta_i) \leqslant p\pi\right] = p,$$

so that since we measure $\mu(O_i) = k$, we have the following result.

Proposition 15 *In a $N \times N$ image, the number of false alarms associated to a segment O_i containing k aligned points (up to precision p) among l is*

$$NFA(O_i) = N^4 \sum_{k'=k}^{l} \binom{n}{k'} p^{k'} (1-p)^{l-k'}. \tag{7.6}$$

7.3.2 Thresholds

The tail of the Binomial law,

$$P(k,l) = \sum_{k'=k}^{l} \binom{n}{k'} p^{k'} (1-p)^{l-k'}, \tag{7.7}$$

involved in (7.6), is not easy to handle mathematically. We now give a more practical estimate, that yields a more simple formula for the thresholds $k(l, p, \varepsilon)$ defining ε-meaningful segments according to Proposition 15 and Definition 19.

Theorem 13 (Hoeffding's inequality) *Let k, l be integers with $0 \leqslant k \leqslant l$, and p a real number such that $0 < p < 1$, Then if $r = k/l \geqslant p$, we have the inequalities*

$$P(k,l) \leqslant \exp\left(lr \ln \frac{p}{r} + l(1-r) \ln \frac{1-p}{1-r} \right) \leqslant \exp(-l(r-p)^2 h(p)) \leqslant \exp(-2l(r-p)^2), \tag{7.8}$$

where h is the function defined on $]0, 1[$ by

$$h(p) = \begin{cases} \dfrac{1}{1-2p} \ln \dfrac{1-p}{p} & \text{if} \quad 0 < p < \frac{1}{2}, \\[3mm] \dfrac{1}{2p(1-p)} & \text{if} \quad \frac{1}{2} \leqslant p < 1. \end{cases} \tag{7.9}$$

Using this theorem, we deduce a sufficient condition for a segment to be meaningful.

Proposition 16 *Assume that $p < 1/2$. In a $N \times N$ image, a segment O_i containing k aligned points (up to precision p) among l is ε-meaningful as soon as*

$$k \geqslant pl + \sqrt{\frac{4 \ln N - \ln \varepsilon}{h(p)}} \sqrt{l}, \tag{7.10}$$

where $h(p)$ is given by (7.9).

Proof :

Assume that (7.10) is satisfied. If we set $r = k/l$, then $r \geqslant p$ and

$$l(r-p)^2 \geqslant \frac{4 \ln N - \ln \varepsilon}{h(p)}.$$

From (7.8) we deduce that

$$P(k,l) \leqslant \exp(-l(r-p)^2 h(p)) \leqslant \exp(-4 \ln N + \ln \varepsilon) = \frac{\varepsilon}{N^4},$$

which means, by definition, that the segment S is ε-meaningful. $\qquad\square$

7.3.3 Maximal alignments

When a segment O_i is meaningful, it generally raises several other meaningful segments, obtained by extending or reducing a little bit the length of O_i. This phenomenon may be avoided by limiting the search to maximal segments, that is segments whose associated number of false alarms is minimal with respect to inclusion.

Definition 20 *A segment O_i is maximal if*

- *it does not contain a strictly more meaningful segment :*

$$\forall j \neq i, \ O_j \subset O_i, \quad NFA(O_j) \geqslant NFA(O_i);$$

- *it is not contained in a more meaningful segment :*

$$\forall j \neq i, \ O_j \supset O_i, \quad NFA(O_j) > NFA(O_i).$$

Now, an important question for maximal segments (and maximal structures in general) is to know whether they can intersect or not. We shall see that even if no complete proof exists, it seems very likely that two distinct maximal segments on the same line cannot meet.

Conjecture 1 ([9]) *If $(l, l', l'') \in \{1.. + \infty\}^3$ and $(k, k', k'') \in \{0..l\} \times \{0..l'\} \times \{0..l''\}$, then*

$$\min\left(p, P(k,l), P(k + k' + k'', l + l' + l'')\right) < \max\left(P(k + k', l + l'), P(k + k'', l + l'')\right). \quad (7.11)$$

This conjecture passed many numerical tests [9], and an asymptotic proof (when $l \to +\infty$) exists. Moreover, Conjecture 1 has been proved to be a consequence of the following conjecture : the function

$$(k, l) \ \mapsto \ \frac{\displaystyle\int_0^p x^{k-1}(1-x)^{l-k}dx}{\displaystyle\int_0^1 x^{k-1}(1-x)^{l-k}dx}$$

has negative curvature on the domain $D_p = \{(k, l) \in \mathbb{R}_+^2, \ p(l-1) + 1 \leqslant k \leqslant l\}$.

Theorem 14 *Suppose that Conjecture 1 is true. Then, any two maximal meaningful segments lying on the same straight line have no intersection.*

Proof :

Suppose that one can find two maximal segments $(k + k', l + l')$ and $(k + k'', l + l'')$ that have a non-empty intersection (k, l) (the couple (k, l) we attach to each segment represents the number of aligned points (k) and the segment length (l)). Then, according to Conjecture 1 we have

$$\min\left(p, P(k,l), P(k + k' + k'', l + l' + l'')\right) < \max\left(P(k + k', l + l'), P(k + k'', l + l'')\right).$$

If the left hand term is equal to p, then we have a contradiction since one of $(k + k', l + l')$ or $(k + k'', l + l'')$ is strictly less meaningful than the segment $(1, 1)$ it contains. If not, we have another contradiction because one of $(k + k', l + l')$ or $(k + k'', l + l'')$ is strictly less meaningful than one of (k, l) or $(k + k' + k'', l + l' + l'')$. $\qquad \square$

7.4 Detection of contrasted edges

7.4.1 Model

Let us now consider a classical example in Computer Vision : the detection of contrasted edges. More precisely, we would like to find out, among the level lines of a discrete image $u : \Omega \to \mathbb{R}$, the curves along which the image gradient remains large enough to be meaningful. Let us note $O_1, O_2, \ldots O_M$ the curves to be investigated. On each curve O_i, we compute

$$\mu(O_i) = \min_{\mathbf{x} \in O_i} |Du(\mathbf{x})|, \tag{7.12}$$

where $Du(\mathbf{x})$ is a local approximation of the image gradient at point $\mathbf{x}$. Now, how to apply Helmholtz Principle ? We could assume that the random variables $|Du(\mathbf{x})|, \mathbf{x} \in \Omega$ are independent, but assuming a uniform distribution is questionable. Indeed, there is no uniform distribution in $\mathbb{R}^+$ (so we would have to choose an arbitrary maximum value), and it is well known that the distribution of $|Du(\mathbf{x})|$ in a natural image is concentrated around small values. Hence, it is more natural to use an a contrario model based only on the independence of the values of $|Du(\mathbf{x})|$, while taking the empirical distribution of $|Du(\mathbf{x})|$ as an a priori distribution. If we define the repartition function of $|Du|$ by

$$H(\mu) = \frac{|\{\mathbf{x} \in \Omega, \ |Du(\mathbf{x})| \geqslant \mu\}|}{|\Omega|}, \tag{7.13}$$

our model will be that

$$\mathrm{P}\left[|D\tilde{u}(\mathbf{x})| \geqslant \mu\right] = H(\mu).$$

Now, if $l(O_i)$ is the maximum number of independent points on O_i, then we have

$$\mathrm{P}\left[\mu(\tilde{O}_i) \geqslant \mu\right] = H(\mu)^{l(O_i)}.$$

This yields the following result.

Proposition 17 *The number of false alarms associated to the curve O_i with minimal contrast $\mu(O_i)$ is*

$$NFA(O_i) = M \cdot H\big(\mu(O_i)\big)^{l(O_i)}, \tag{7.14}$$

where $l(O_i)$ is the maximum number of independent points on O_i and H is the repartition function of $|Du|$ defined by (7.13).

7.4.2 Numerical issues

Numerically, the model above may be realized differently depending on the choice of the interpolation method.

- A first solution consists in using nearest neighbor interpolation for the computation of the level lines, yielding discrete curves made of vertical and horizontal segments. In that case, $|Du|$ may be estimated by $|u(x+1,y) - u(x,y)|$ for a vertical segment, by $|u(x,y+1) - u(x,y)|$ for an horizontal segment, and the number of independent points on a curve is approximately half the number of vertical plus horizontal segments. A 2×2 finite difference scheme may be used instead to estimate $|Du|$ less locally. In this case, $|Du|$ will be estimated at the junctions between segments, and the length of the curve will be smaller, since two successive points $\mathbf{x}, \mathbf{x}' \in \mathbb{Z}^2$ must satisfy $|\mathbf{x} - \mathbf{x}'| \geqslant 2$ to be independent.

- A more accurate solution consists in using bilinear interpolation, both for the computation of the level lines and for the estimation of $|Du|$. The number of independent points on a curve can then be approximated by $l/3$, where l is the Euclidean length of the curve (note that this is rather an upper bound).

Now we have to precise which curves will be used for boundary detection. We have two possibilities.

7.4.3   Meaningful boundaries (closed)

We may look only for closed curves, yielding what we shall call *meaningful boundaries*. In practice, $(O_i)_i$ will be the set of λ-level lines of u, where λ belong to a fixed grid $q\mathbb{Z}$. The number of false alarms is given by (7.14), where M is the total number of $q\mathbb{Z}$-level lines of u. A notion of maximal meaningful boundaries may then be defined with respect with the inclusion tree of level lines.

7.4.4   Meaningful edges

If we want to find non-closed curves (*meaningful edges*), we can choose for $(O_i)_i$ the set of all pieces of level lines. In practice, on each level line we choose a maximal set $\mathbf{x}_1, \mathbf{x}_2, \ldots \mathbf{x}_n$ of independent points and then consider only the $n(n-1)/2$ sub-curves whose endpoints belong to this set. The NFA computation is the same as for meaningful boundaries, except that the number of objects (M) in (7.14) is much larger. If maximal edges are defined like in Definition 20, one can show easily that maximal edges are disjoint.

References

[1] L. Alvarez, F. Guichard, P.L. Lions, J.M. Morel, "Axioms and fundamental equations of image processing", *Archives for Rational Mechanics* 123, pp. 199-257, 1993.

[2] G. Aronsson, "Extensions of functions satisfying lipschitz conditions", *Arkiv för Matematik* 6, 1967.

[3] M. Ayer, H.D. Brunk, G.M. Ewing, W.T. Reid, E. Silverman, "An empirical distribution function for sampling with incomplete information", *Ann. Math. Statist.* 26, pp. 641-647, 1955.

[4] A. Canny, "A computational approach to edge detection", *IEEE Trans. on PAMI* 8:6, pp. 679-698, 1986.

[5] V. Caselles, B. Coll et J.-M. Morel, "A Kanizsa programme", *Progress in Nonlinear Differential Equations and their Applications*, 25, pp. 35-55, 1996.

[6] V. Caselles, J-M. Morel, C. Sbert, "An Axiomatic Approach to Image Interpolation", *IEEE Transactions On Image Processing* 7:3, pp. 376-379, 1998.

[7] M.G. Crandall, H. Ishii, P.-L. Lions, "User's guide to viscosity solutions of second order partial differential equations", *Bulletin of Amer. Math. Society* 27, pp. 1-67, 1992.

[8] M.G. Crandall, H. Ishii, P.-L. Lions, "Viscosity solutions : a primer", in *Viscosity Solutions and Applications*, Springer Lecture Notes in Mathematics 1660, 1997.

[9] A. Desolneux, L. Moisan and J.-M. Morel, "Meaningful Alignments", *International Journal of Computer Vision*, 40:1, pp.7-23, 2000.

[10] A. Desolneux, L. Moisan and J.-M. Morel, "Edge Detection by Helmholtz Principle", *Journal of Mathematical Imaging and Vision*, 14:3, pp. 271-284, 2001.

[11] A. Desolneux, L. Moisan and J.-M. Morel, "Maximal Meaningful Events and Applications to Image Analysis", *Annals of Statistics*, 31:6, 2003.

[12] M.D. Fairchild, *Color Apperance Models*, Addison Wesley, 1998.

[13] E. Giusti, *Minimal Surfaces and functions of bounded variation*, Birkhäuser, 1984.

[14] F. Guichard, J.-M. Morel, *Image Analysis and PDE's*, to appear.

[15] R.M. Haralick, S.R. Sternberg, X. Zhuang, "Grayscale Morphology", *Comp. Vision, Graphics, Image Proc.* 35, 1986.

[16] R. Jensen, "Uniqueness of Lipschitz extension : minimizing the sup norm of the gradient", *Arch. Rat. Mech. Anal.* 123, pp. 51-74, 1993.

[17] G.Kanizsa, *Grammatica del Vedere*, Il Mulino, Bologna, 1980. Traduction : *La grammaire du voir*, Diderot Editeur, Arts et Sciences, 1996.

[18] Yong Hoon Lee, Saleem A. Kassam, "Generalized Median Filtering and Related Nonlinear Filtering Techniques".

[19] D. Lowe, *Perceptual Organization and Visual Recognition*, Kluwer Academic Publishers, 1985.

[20] G. Matheron, *Random Sets and Integral Geometry*, John Wiley, New York, 1975.

[21] P. Maragos, "Tutorial on advances in morphological image processing and analysis", *Optical Engineering* 26:7, pp. 623-632, 1987.

[22] D. Marr, *Vision*, Freeman and Co., 1982.

[23] L. Moisan, "Affine Plane Curve Evolution: a Fully Consistent Scheme", *IEEE Transactions On Image Processing*, vol. 7:3, pp. 411-420, 1998.

[24] P. Monasse, F. Guichard, "Fast Computation of a Contrast Invariant Image Representation" *IEEE Transactions on Image Processing* 9:5, pp. 860-872, 2000.

[25] T. Robertson, F.T. Wright, R.L. Dykstra, *Order Restricted Inference*, John Wiley and Sons, 1988.

[26] L. Rudin, S. Osher, C. Fatemi, "Nonlinear total variation based noise removal algorithm", *Physica D* 60, pp. 259-268, 1992.

[27] L. Rudin, S. Osher, "Total Variation based image restoration with free local constraints", *Proc. IEEE Int. Conf. on Image Processing* 1, Austin, pp. 31-35, 1994.

[28] G. Sapiro, A. Tannenbaum, "On affine plane curve evolution", *J. of Functional Analysis* 119, pp. 79-120, 1994.

[29] J. Serra, *Image analysis and mathematical morphology*, Academic Press, 1982.

[30] M. Sigman M, G.A. Cecchi, C.D. Gilbert, and M.O. Magnasco. "On a common circle: Natural scenes and Gestalt rules", *Proc. Nat. Acad. Sci. USA* 98:4, pp. 1935-1940, 2001.

[31] C. V. Stewart, "MINPRAN : a new robust estimator for Computer Vision", *IEEE Trans. on Pattern Analysis and Machine Intelligence* 17, pp. 925-938, 1995.

[32] A.N. Tikhonov, "Regularization of incorrectly posed problems", *Soviet Mathematics Doklady* 4, pp. 1624-1627, 1963.

[33] M. Unser, A. Aldroubi, M. Eden, "Fast B-Spline Transforms for Continuous Image Representation and Interpolation", *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13:3, pp. 277-285, 1991.

[34] C.R. Vogel, M.E. Oman, Iterative methods for total variation denoising, *SIAM J. Sci. Comput.* 17:1, pp. 227-238, 1996.

[35] M. Wertheimer, "Untersuchungen zur Lehre der Gestalt", II, *Psychologische Forschung* 4, pp. 301-350, 1923.

[36] A.P. Witkin, "Scale-Space filtering", *Proceedings of IJCAI*, Karlsruhe, pp. 1019-1021, 1983.

[37] W.P. Ziemer, *Weakly differentiable functions*, Springer-Verlag, New-York, 1989.