Art & Perception

Boundaries, Transitions and Passages

Jan Koenderink[1,2], Andrea van Doorn[1,2], Baingio Pinna[3] and Johan Wagemans[1]

[1] University of Leuven (KU Leuven), Laboratory of Experimental Psychology,
Tiensestraat 102, Box 3711, 3000 Leuven, Belgium
[2] Experimental Psychology, Utrecht University,
Heidelberglaan 1, 3584 CS Utrecht, The Netherlands
[3] Department of Humanities and Social Sciences, University of Sassari,
Via Roma 151, 07100 Sassari, Italy

Supplementary Material

Supplement: The formalism of transitions

A point of size s located at $\{x_0, y_0\}$ is a linear operator with point spread function

$$p(x,y,x_0,y_0,s) = \frac{\mathrm{e}^{-\frac{(x-x_0)^2+(y-y_0)^2}{2s^2}}}{2\pi s^2}, \tag{1}$$

which is normalised to unit overall weight:

$$\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} \frac{\mathrm{e}^{-\frac{(x-x_0)^2+(y-y_0)^2}{2s^2}}}{2\pi s^2}\,\mathrm{d}x\,\mathrm{d}y = 1. \tag{2}$$

Here we let "p" stand for "point operator (see figure S1).

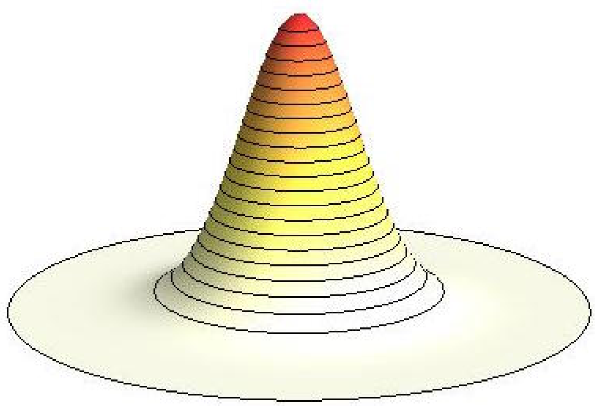

Figure S1. The receptive field profile of the point operator. It is positive throughout, although the weight falls off steeply with distance to the center. Here we use a temperature scale, with red for the most positive value, negative values becoming blue (not present here), values close to zero whitish.

A point samples the image $I(x,y)$, that is to say

$$\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} p(x,y,x_0,y_0,s)I(x,y)\,\mathrm{d}x\,\mathrm{d}y = I(x_0,y_0,s), \tag{3}$$

thus a point "takes a bite out of the image" so to speak. The value obtained is the image intensity at that location as seen by the acuteness of the point. That is its operational definition. In order to simplify notation we will often assume the planar coordinates $\{x, y\}$ as "understood", and write $I(s)$ for $I(x, y, s)$. The scale s has to be indicated in all cases, except for the original image, which will be denoted as simply I.

Thus

$$p(s) \circ I = I(s), \tag{4}$$

where $I(s)$ is the image at level of resolution s. Here the operator "$\circ$" denotes convolution. Thus one might say $I(s)$ is a "Gaussian blurred copy" of the image. However, from a conceptual perspective "the" image I, that is the image at infinite resolution, does not exist. One may only know images at finite resolution. Because formally $I(s_1) \circ p(s_2) = I(\sqrt{s_1^2 + s_2^2})$, this is not a problem.

Notice that

$$\Delta p = \frac{1}{s} \frac{\partial p}{\partial s}, \tag{5}$$

which implies

$$\Delta I(s) = \frac{1}{s} \frac{\partial I(s)}{\partial s}. \tag{6}$$

This defines the "scale-space" $I(s)$ of the (ideal!) image I, where I stands formally for $I(0)$. The partial differential equation is the diffusion equation, of which the point is a kernel.

The negative directional derivative in the x-direction of the point is

$$b(s) = -\frac{\partial p(s)}{\partial x}, \tag{7}$$

or, in expanded notation

$$b(x, y, x_0, y_0, s) = \frac{(x - x_0)\, \mathrm{e}^{-\frac{(x-x_0)^2 + (y-y_0)^2}{2s^2}}}{2\pi s^4}. \tag{8}$$

This is usually called an "edge finder" (see figure S2), which is merely an unfortunate term for a Gaussian derivative operator. Here we let "b" stand for "border operator".

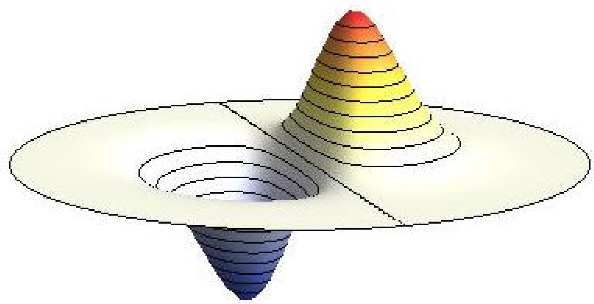

Figure S2. This is a directional derivative of a point, often called "edge finder", in our terminology the "border operator". Formally it behaves like a tangent vector of the Euclidean plane at the given scale. It has positive and negative lobes that exactly balance each other, thus the operator is blind to uniform patches. This receptive field profile may double as the "border brush icon" in our model, they look exactly the same.

The border operator doubles as the "brush operator". If we "paint the border operator" with the "brush operator" we obtain

$$b(s) \circ b(s) = \frac{\partial^2 p(s\sqrt{2})}{\partial x^2} = \ell(s\sqrt{2}), \tag{9}$$

where "ℓ" is known as a formal expression for the Hubel and Wiesel "line finder" (see figure S3). One might call ℓ the "local borderness" in the x-direction. The scale has slightly increased from s to $s\sqrt{2}$.

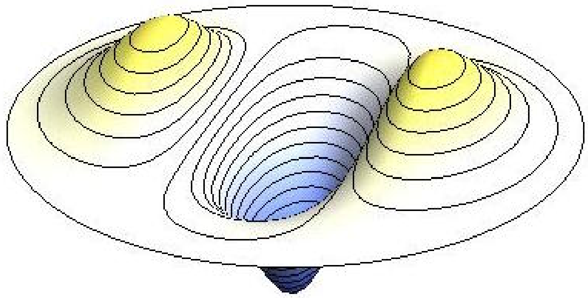

Figure S3. This is the second order directional derivative of a point. It is like the type of simple cell called "line finder" in neurophysiology. The second order derivative equals the concatenation of two first order derivatives, this is the convolution of an edge operator with itself. In our model it will be "borderness painted by the border-icon", that is to say, a local border.

One may have line finders in *any* direction. Summing over all directions removes the directional dependence. By substituting $x \to \varrho \cos\varphi$, $y \to \varrho \sin\varphi$ and integrating over the angle φ, one obtains

$$\frac{1}{2\pi}\int_{-\pi}^{+\pi} \ell(s)\,\mathrm{d}\varphi = \frac{1}{2}\Delta p(s), \tag{10}$$

where we have suppressed the explicit ϱ parameter in accordance with the short-hand notation. This sum is the local overall borderness, irrespective direction.

The Laplacean operator is the difference of two point operators of slightly different scales. Thus it is similar to what the neurophysiologist calls a "DOG ("difference of Gaussians") receptive field (see figure S4). This interpretation also indicates that the Laplacean is proportional to the derivative with respect to scale.

Since the Laplacean of the point operator is proportional to the derivative with respect to scale,one has

$$\int_{0}^{+\infty} s\,\Delta I(s)\,\mathrm{d}s = 0 \quad \text{outside the origin,} \tag{11}$$

whereas

$$\int_{s_0}^{+\infty} s\,\Delta I(s)\,\mathrm{d}s = -\frac{1}{2\pi s_0^2} \quad \text{at the origin,} \tag{12}$$

implying that the integral over the full scale domain is the Dirac delta function.

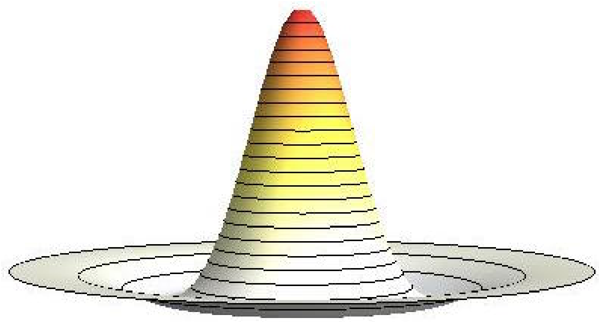

Figure S4. This is the difference of a point and another point $\sqrt{2}$–times as large. It is virtually indistinguishable from the scale derivative of a point, or the Laplacean operator. The positive center has a negative surround. Integrated over all space one gets zero, thus such operators are blind to uniform patches.

Putting things together we have that s times the overall borderness at scale s equals the scale derivative at scale s. What holds locally *eo ipso* hold in the image scale space because convolution is associative and commutative. Thus *the borderness distribution of an image is proportional to its derivative with respect to scale.*

Inverting this yields the theorem that adding the borderness over all scales is the same as integrating the derivative with respect to scale over scale, that is to say, simply the image again. This is a major insight:

> *the image is the sum of the borderness over all scales.*

Thus one may analyse the image into transition areas and synthesise it again. So where is the gain—since we started with an image and ended with it? The gain is that after the analysis one may selectively handle the components in the synthesis. Thus one may "filter" based on borderness.

Because the "ideal" image is a non-entity, one may want to split the integral over scale into three parts. At the lowest resolutions one is not interested in border regions at all. One may simply use the images in its most blurred form as a given basis. It will be almost uniform, so it may also be ignored altogether. At the highest resolutions one lacks data. Infinite resolution is not to be had. The solution is simply to omit the integration over the highest resolutions. Thus one is left with an integration over a slice of scales that is indeed availabe. One ends up with a very practical algorithm.