


B. The concept of convolution

Out[100]=



From the book:

ter Haar Romeny, B. M.:

**Front-end vision and multi-scale image analysis:
Multi-scale computer vision theory
and applications, written in Mathematica.**

ISBN: 978-1402015038

Publisher: Springer Dordrecht (2003)

Download: www.romeny.info/FEV

Initialization

```
In[7]:= url = "https://www.romeny.info/FEV-CD/images/";  
SetOptions[ArrayPlot, ColorFunction -> GrayLevel];  
$HistoryLength = 0;  
gD[im_, nx_, ny_,  $\sigma$ ] := DerivativeFilter[im, {ny, nx},  $\sigma$ ];  
Off[DerivativeFilter::arg1];  
AutoCollapse[] := (If[$FrontEnd != $Failed,  
  SelectionMove[EvaluationNotebook[], All, GeneratedCell];  
  FrontEndTokenExecute["SelectionCloseUnselectedCells"]]);
```

B.1 Convolution

```
In[13]:= Show[Import[url <> "retinapsf.gif"], ImageSize -> 300]
```

```
Out[13]=
```

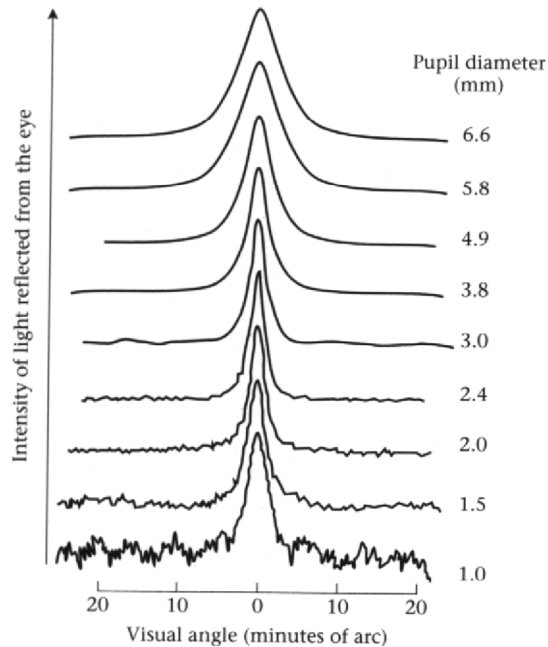


Figure B.1 Experimental measurements of light that has been reflected from a human eye looking at a fine line ('line spread function'). The reflected light has been blurred by double passage through the optics of the eye. The amount of unsharpness is a function of the pupil diameter: with wider pupil opening the line got more blurred. Source: Campbell and Gubish, 1996. Taken from [Wandell1995].

We take the example of the not-exactly-sharp projection of images on the human retina. Lens errors, cornea distortions, the retinal nerve tissue: they all contribute to the 'smearing effect'.

In fact, every point in the stimulus gives rise to the same blurring effect. Figure 1 shows some actual measurements at various pupil diameters. Suppose that the profile of the blur function is a Gaussian kernel. Then every point of the stimulus will become a Gaussian function when projected on the retina. We study the 1-D case, and call the blur function $\mathbf{g}[\mathbf{x}, \sigma]$. The blur function is also called the *filter* or the *kernel*, or the *operator* (sometime we encounter the name *stencil* or *template*).

```
In[14]:= gauss[x_, sigma_] := 
$$\frac{1}{\sqrt{2\pi}\sigma^2} \text{Exp}\left[-\frac{x^2}{2\sigma^2}\right];$$

```

We define the input signal as $\mathbf{f}[\mathbf{x}]$, and the output signal as $\mathbf{g}[\mathbf{x}]$. We can think of the input signal as a series of points next to each other, and each of these input points gives rise to a blurred output point. We plot the input point function $\mathbf{f}[\mathbf{x}]$ (a very narrow function around the point $\mathbf{x}=0$) at positions $\mathbf{x}=0$ and $\mathbf{x}=2$:

```

In[15]:= fpnt[x_] = If[-0.05 < x < 0.05, 1, 0];
p1 = Plot[fpnt[x], {x, -3, 3},
  PlotRange -> {0, 1}, PlotStyle -> {Red, Thick}];
p2 = Plot[fpnt[x - 2], {x, -3, 3}, PlotRange -> {0, 1},
  PlotStyle -> {Red, Thick}];
GraphicsRow[{p1, p2}, ImageSize -> 300]

```




Figure B.2 The sampling function, at two arbitrary positions.

As an input signal we take an example function consisting of a sum of sine functions of different frequencies. We can approximate this function as a set of points very close to each other. The amplitude of each point is modulated with the input function, so we plot the product $f[\alpha] \text{fpnt}[x-\alpha]$.

```

In[18]:= f[x_] = Sin[x] + 0.5 Sin[5 x];
p1 = Plot[f[x], {x, -4, 4}];
p2 = Show[Table[Plot[f[α] × fpnt[x - α], {x, -4, 4},
  PlotPoints -> 160, PlotRange -> {-1.5, 1.5}], {α, -4, 4, .1}]];
Show[GraphicsRow[{p1, p2}], ImageSize -> 500]

```




Figure B.3 Left: The input signal $f(x)$, right: the sampled representation.

Every point is blurred on its own, so we replace every pointfunction $\text{fpnt}[\alpha]$ by its blurred version $g[\alpha]$. Every point is widened substantially, and the final step is adding all these responses together.

```

In[22]:= p1 = Show[
  Table[Plot[f[α] × gauss[x - α, 1], {x, -4, 4}, PlotRange → {-0.75, 0.75},
    PlotPoints → 160, PlotStyle → Thickness[0.003]], {α, -4, 4, .1}];
  h[x_] = Sum[f[α] × gauss[x - α, 1], {α, -6, 6, .1}];
  p2 = Plot[h[x], {x, -4, 4}];
  Show[GraphicsRow[{p1, p2}], ImageSize → 500]

```

Out[25]=




Figure B.4 Left: every sample gives rise to a kernel, at the respective position with the amplitude of the sample. Right: the summed response is the convolution.

In the limit of making the plot functions smaller and smaller we get a summation described by an integral: the convolution-integral. So we may write for the output $\mathbf{h}[\mathbf{x}]$:

$$\mathbf{h}[\mathbf{x}] = \int_{-\infty}^{\infty} \mathbf{f}[\alpha] \times \mathbf{g}[\mathbf{x} - \alpha, \sigma] \, d\alpha$$

It describes exactly how the output $\mathbf{h}[\mathbf{x}]$ looks like when an input signal $\mathbf{f}[\mathbf{x}]$ is processed by a system $\mathbf{g}[\mathbf{x}]$. We say that the signal $\mathbf{f}[\mathbf{x}]$ is filtered with a filter or kernel $\mathbf{g}[\mathbf{x}]$. We also call $\mathbf{g}[\mathbf{x}]$ the point spread-function of the system. Note that the convolution-integral integrates from $-\infty$ to ∞ , indicating that we integrate over the whole length of the signal. The parameter α is the so-called shift variable. (sometimes the convolution integral is called the shift-integral).

Actually, the sum above is an approximation. The exact output is given by the convolution-integral. Let us plot the integral, and draw your conclusions yourself. We have approximated the analytical solution quite well:

```

In[26]:= h[x] = Integrate[f[α] × gauss[x - α, 1] dα

```

Out[26]=

$$0.606531 \text{ Sin}[x] + 1.86333 \times 10^{-6} \text{ Sin}[5 x]$$

```
In[27]:= Plot[Evaluate[h[x]], {x, -4, 4}, ImageSize -> 160]
Out[27]=
```




Figure B.5 The convolution as an exact integral of the analytical input function and kernel.

The output signal looks different from the input signal, and this is exactly what we wanted: we have filtered the high frequencies out of the signal. The frequency $\text{Sin}[5x]$ is virtually eliminated, as the figure above is almost the low frequency $\text{Sin}[x]$ function, as we put in the input. This low frequency filter is however 40% attenuated due to the filtering (check the amplitudes).

A short form for writing the convolution integral is the symbol \otimes , the convolution operator. The function h is the convolution of the function f with g

$$h[x] = f[x] \otimes g[x]$$

The convolution operator is a *linear* operator:

$$\begin{aligned} (h_1[x] + h_2[x]) \otimes g[x] &= h_1[x] \otimes g[x] + h_2[x] \otimes g[x]; \\ (a f[x]) \otimes g[x] &= a (f[x] \otimes g[x]) \end{aligned}$$

In 2-D we get:

$$h[x_, y_] := \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f[\alpha, \beta] \times g[x - \alpha, y - \beta] \, d\alpha \, d\beta$$

B.2 Convolution is a product in the Fourier domain

It is often not so easy to calculate the convolution integral. It is often much easier to calculate the integral of the convolution in the Fourier domain. This section explains how this works. We start from the convolution of the function $h(x)$ with the kernel $g(x)$:

$$f(x) = \int_{-\infty}^{\infty} h(y) g(x - y) \, dy$$

The Fourier transform $F(\omega)$ of $f(x)$ is then by definition

$$F(\omega) = \int_{-\infty}^{\infty} f(x) e^{-i\omega x} \, dx = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} h(y) g(x - y) \, dy e^{-i\omega x} \, dx$$

We reorder the terms in the integral and make the substitution $x - y = \tau$, so we have $x = \tau + y$ and $e^{-i\omega x} = e^{-i\omega\tau} \cdot e^{-i\omega y}$:

$$\int_{-\infty}^{\infty} h(y) \int_{-\infty}^{\infty} e^{-i\omega x} g(x - y) dx dy =$$

$$\int_{-\infty}^{\infty} h(y) \int_{-\infty}^{\infty} e^{-i\omega(\tau+y)} g(\tau) d(\tau + y) dy$$

When we integrate to τ , we keep y constant, so $d(\tau+y)$ is equal to $d\tau$. So we get:

$$\int_{-\infty}^{\infty} h(y) \int_{-\infty}^{\infty} e^{-i\omega\tau} e^{-i\omega y} g(\tau) d\tau dy$$

Because $e^{-i\omega y}$ is constant in the integration to τ , we may bring it as a constant outside the integral:

$$\int_{-\infty}^{\infty} h(y) e^{-i\omega y} dy \int_{-\infty}^{\infty} e^{-i\omega\tau} g(\tau) d\tau = H(\omega) \cdot G(\omega)$$

where $H(\omega)$ is the Fourier transform of the function $h(x)$ and $G(\omega)$ is the Fourier transform of the kernel $g(x)$. So the Fourier transform of a convolution is the product of the Fourier transform of the filter with the Fourier transform of the signal. To put it otherwise:

$$f(x) = h(x) \otimes g(x)$$

$$\updownarrow \quad \updownarrow \quad \updownarrow$$

$$F(\omega) = H(\omega) \cdot G(\omega)$$

We can calculate a convolution $h \otimes g$ by calculating the Fourier transforms $H(\omega)$ and $G(\omega)$, multiply them to get $F(\omega)$, and we get $f(x)$ by taking the inverse Fourier transform of $F(\omega)$. Often this is a much faster method than calculating the convolution integral, because the routines that calculate the Fourier transform are so fast. We look at an example where we filter noise from the data. We simulate a signal of a noisy ECG recording by some sine functions:

```
In[52]:= data = Table[
  N[Sin[ $\frac{2\pi}{256} 2 t$ ] + Sin[ $\frac{2\pi}{256} 6 t$ ] + 0.75 (Random[ ] - 1/2)], {t, 256}];
ListPlot[data, AspectRatio -> .3, Joined -> True,
  ImageSize -> 400, AxesLabel -> {"time", "Ampl"}]
```




Figure B.6 A discrete input signal with additive uncorrelated uniform noise.

We make a filter with a Gaussian shape and shift it to the origin:

```
In[30]:=  $\sigma = 5.$ ; kernel = Table[gauss[x,  $\sigma$ ], {x, -128, 127}];
kernel = RotateLeft[kernel, 128];
ListPlot[kernel, PlotRange -> {{0, 256}, All}, Joined -> True,
  AspectRatio -> .2, AxesLabel -> {"Freq", "Ampl"}, ImageSize -> 400]
```




Figure B.7 The Gaussian kernel in the spatial domain, shifted to the origin. The right half of the frequency axis (i.e. 129-256) is often called the negative frequency axis.

This is a low-pass filter. We now filter (convolve) in the Fourier domain (note that a space denotes multiplication in *Mathematica*):

```
In[50]:= conv = Sqrt[256] Chop@InverseFourier[Fourier[data]  $\times$  Fourier[kernel]];
ListPlot[conv, AspectRatio -> .3, Joined -> True,
  AxesLabel -> {"time", "Ampl"}, ImageSize -> 400]
```




Figure B.8 The filtered result.

As expected, the noise did not pass the filter.