Linear Filtering and Basis Functions (I)

We put these models into the context of the literature on linear filtering and fourier analysis. This is an advanced section which gives greater understanding but is not required for a basic introduction.

As discussed earlier, simple cell models apply *linear filters* to images and cells at different spatial locations perform *convolution* $*$ by applying the same filter $\vec{w}$ across the image:

$$S(\vec{x}) = \vec{w} * I(\vec{x}) = \sum_{\vec{y}} w(\vec{x} - \vec{y})I(\vec{y}).$$

It is also convenient to approximate this (take the continuum limit) and express this as an integral:

$$S(\vec{x}) = \int_{\vec{y}} w(\vec{x} - \vec{y})I(\vec{y})d\vec{y}.$$

This continuum limit is a good approximation, if the summation $\sum_{\vec{y}}$ is over a dense set of positions $\vec{y}$, and enable certain type of analysis (e.g., showing that a center-surround cell model sums, approximately, to zero).

Convolving by a Gaussian and Derivatives of a Gaussian

Convolving an image by a linear filter produces an output image $S(\vec{x})$ whose form depends on the type of filter $\vec{w}$. For example, if $w(\vec{x})$ is a Gaussian function $G(\vec{x}; \sigma) = \frac{1}{2\pi\sigma^2} \exp\{-(x_1^2 + x_2^2)/(2\sigma^2)\}$ then convolution effectively just smooths the image by taking a linear weighted average. If $\vec{w}$ is a derivative of the Gaussian in the x_1 direction, $w(\vec{x}) = \frac{d}{dx_1} G(\vec{x}; \sigma)$, then this filter gives a large response to *edges*, positions $\vec{y}$ where the intensity $I(\vec{y})$ changes abruptly, and has small responses in places where the image intensity changes slowly.

We can better understand images, and linear filtering, by using *functional analysis*. This states that an image, or any signal, can be expressed uniquely as a weighted sum of *basis functions*:

$$I(\vec{x}) = \sum_i \alpha_i b_i(\vec{x}), \tag{1}$$

where the $b_i(\vec{x})$ are *basis functions* and the $\{\alpha_i\}$ are *coefficients*. These basis functions are usually chosen to be *orthonormal*, so that $\sum_{\vec{x}} b_i(\vec{x}) b_j(\vec{x}) = \delta_{ij}$ (= 1 if $i = j$ and = 0 if $i \neq j$). If the basis functions are orthogonal then the coefficients α can be obtained by:

$$\alpha_i = \sum_{\vec{x}} I(\vec{x}) b_i(\vec{x}). \tag{2}$$

The *principle of superposition* states that we can determine the output S as a weighted combination of the outputs of the basis functions:

$$S(\vec{x}) = \sum_i \alpha_i S_i(\vec{x}), \quad \text{where } S_i(\vec{x}) = \sum_{\vec{y}} w(\vec{x} - \vec{y}) b_i(\vec{y}). \qquad (3)$$

This implies that if we know the response $S_i(.)$ to each basis function $b_i(.)$, then we can predict the response to any input. This is an attractive property which, if it holds, enables us to measure the receptive field of a linear neuron, or a thresholded linear neuron, from a limited set of stimuli.

Fourier analysis deals with a special class of basis functions. These are sinusoids, i.e. of form $\sin \omega x, \cos \omega x$. Then the α's are the *fourier transform* of the image. If we restrict ourselves to an image defined on a lattice (i.e. so that x_1, x_2 each take a finite number of values, as on a digital camera) then this is the *discrete fourier transform*. But if we allow x_1, x_2 to take continuous values, then we get the *fourier transform*:

$$I(\vec{x}) = \frac{1}{2\pi} \int \hat{I}(\vec{\omega}) \exp\{-i\vec{\omega} \cdot \vec{x}\} d\vec{\omega} \tag{4}$$

$$\hat{I}(\vec{\omega}) = \frac{1}{2\pi} \int I(\vec{x}) \exp\{i\vec{\omega} \cdot \vec{x}\} d\vec{x} \tag{5}$$

Here $\exp\{i\vec{\omega} \cdot \vec{x}\} = \cos(\vec{\omega} \cdot \vec{x}) + i \sin(\vec{\omega} \cdot \vec{x})$. Note that if $I(.)$ is symmetric, $I(\vec{x}) = I(-\vec{x})$, then $\hat{I}(\vec{\omega})$ is also symmetric $\hat{I}(-\vec{\omega}) = \hat{I}(\vec{\omega})$. Observe that equations (4,5) correspond to equations (1,2) for special choices of the basis functions (and changing from discrete to continuous $\vec{x}$).

Fourier analysis is particularly important because it gives us a way to represent non-local structure of images in terms of *frequencies* ω. The high frequencies (large $|\vec{\omega}|$) represent image patterns which change rapidly while the lower frequencies (small $|\vec{\omega}|$) represent slowly changing patterns. In particular, if an image pattern is *periodic*, like the stripes on a zebra, then it can be expressed in form:

$$I(\vec{x}) = \sum_n A_n \cos(2\pi n\vec{\omega}_0 \cdot \vec{x}),$$

where $\vec{\omega}_0$ is the basic frequency and n denote integers. Then the Fourier transform is only non-zero at integer multiples of the basic frequency $\vec{\omega} = \vec{\omega}_0$. Hence periodic image patterns, such as *textures*, have very simple descriptions in Fourier space.

If we blur the image, by convolving with a Gaussian $G(\vec{x}; \sigma)$, to obtain $G * I(\vec{x})$, then the high frequencies of the image $\vec{I}$ will be smoothed out. By the *convolution theorem* the fourier transform of $G * I(\vec{x})$ is the product of the fourier transforms of G and $\vec{I}$. The F.T. of a Gaussian is also a Gaussian $\exp\{-|\vec{\omega}|^2(\sigma^2/2)\}$. Hence we can express the convolved image as a weighted combination of sinusoids where the high frequency weights are decreased by $\exp\{-|\vec{\omega}|^2(\sigma^2/2)\}$:

$$\vec{I}(\vec{x}) = \frac{1}{2\pi} \int \hat{I}(\vec{\omega}) \exp\{-i\vec{\omega} \cdot \vec{x}\} \exp\{-|\vec{\omega}|^2(\sigma^2/2)\} d\vec{\omega}.$$

If we increase the blurring, by increasing the variance σ^2, we will make the high frequency coefficients small. Blurring the image can be obtained by defocusing your eyes so that the image is seen out of focus. The receptive fields of cells occurs at a range of different scales, corresponding to convolving with Gaussians of different variances.

The superposition principle, combined with the use of basis functions, shows that we can determine the receptive fields of linear neurons by stimulating them with sinusoids. Sinusoids can be used as basis functions and superposition can be used to predict the response to stimuli which have not been seen yet (i.e. as superpositions of those stimuli to which the response is known). This, however, is rarely done.