

Sparsity, matched filters, and natural images

- ▶ Next, we consider receptive field models from different perspectives. This includes the use of *sparsity* to suggest receptive field properties based on the statistics of natural images as well as the idea of *matched filters*, which revert to an older idea of receptive fields as feature detectors (Lettvin et al., 1959). Sparsity was proposed by Barlow (1961) as a general principle for modeling the brain based on the observation that typically only a small number of neurons are active. It was developed as a way to predict receptive field properties by Olshausen and Field (1996). It is natural to ask whether the receptive fields of cells encode basis functions that somehow capture the typical structure of images and represent it in a form that is suitable for later processing.
- ▶ Our starting point is the idea that images, and particularly local regions of images, can be represented as a linear combination of basis functions $I(\vec{x}) = \sum_i \alpha_i \vec{b}_i(\vec{x})$, as we saw in equation (1).

Sparsity and overcomplete bases

- ▶ Consider an image consisting of regions where the intensity varies spatially smoothly and regions where the intensity consists of a number of bright spots, or *impulses*. The smoothly varying regions of the image can be efficiently represented by Fourier analysis, in the sense that we can approximate the intensity by only a small number of weighted sinusoids
- ▶ By contrast, the impulses are much better represented in terms of a basis of impulse functions. It would be inefficient to represent them in terms of sinusoids.
- ▶ In short, different types of basis functions are suitable for different regions of the image.
- ▶ This suggests a strategy of seeking a representation in terms of an overcomplete set of basis functions, in this case sinusoids and impulse functions, and a criterion that selects an efficient representation so that only a small number of basis functions are activated for each image. This requirement is called ℓ_1 *sparsity*.

ℓ_1 sparsity

- ▶ More formally, we represent an image, or local image region, by:

$$I(\vec{x}) = \sum_{i=1}^N \alpha_i b_i(\vec{x}),$$

where the $\{b_i\}$ are the basis functions and the $\{\alpha_i\}$ are the coefficients.

- ▶ The number N of bases is bigger than the dimension of the image, and hence the bases are *overcomplete*. Overcompleteness implies that there are many ways to represent the image in terms of these basis functions (by different choices of the α 's) and that we need an additional criterion to select the α 's. The ℓ_1 sparsity criterion proposes that we favor representations that make $\sum_{i=1}^N |\alpha_i|$ small, penalize the weights of the basis functions, and encourage most coefficients to be 0.

ℓ_1 sparsity criterion

- ▶ We represent an image \vec{I} by the approximation $\sum_{i=1}^N \hat{\alpha}_i \vec{b}_i$, where the $\{\hat{\alpha}_i\}$ are chosen to minimize the function:

$$E(\alpha) = \sum_{\vec{x}} (I(\vec{x}) - \sum_{i=1}^N \alpha_i b_i(\vec{x}))^2 + \lambda \sum_{i=1}^N |\alpha_i|. \quad (6)$$

- ▶ The first term penalizes the error of the approximation, and the second term, whose strength is weighted by a parameter λ , penalizes the coefficients $\{\alpha_i\}$. The solution $\hat{\alpha} = \arg \min_{\alpha} E(\alpha)$ cannot be specified in closed form, but $E(\alpha)$ is a *convex* function of α , and efficient algorithms exist for minimizing it to estimate $\hat{\alpha}$. The results of these algorithms can, for example, decompose an image into a sum of sinusoids and a sum of impulse functions.

Sparsity and receptive fields (I)

These ideas give an alternative way to think about the receptive fields of cells in V1. First, observe that V1 has far more cells than the retina or the LGN, and so it has enough neural machinery to implement overcomplete bases. Second, overcomplete bases can be designed for specific image structures of interest (e.g., impulse functions or edges), which enables us to start interpreting the image instead of simply representing it. Third, it relates to the observation that cells in V1 fire *sparsely*, which suggests (Barlow, 1961) that they are tuned to specific stimuli and may relate to metabolic processes (firing a neuron takes energy, which needs to be replenished). Hence the idea that the visual cortex seeks to obtain sparse, and hence presumably more easily interpretable, representations has intuitive appeal.

Sparsity and receptive fields (II)

- ▶ Families of Gabor filters give an overcomplete basis, so they do not specify a unique representation of an image. These issues, and the relations of Gabors to wavelets, are discussed in more detail in (Lee, 1996).
- ▶ Sparsity can be used to derive the properties of receptive fields of cells in V1 from natural images (Olshausen, 1996), see figure (18)(Left). Hence instead of hypothesizing models of receptive fields (e.g., Gabor filters), we can try to predict these receptive fields from studying images. These predictions do give some justification for Gabor functions, but they also suggest other receptive field models that have been experimentally observed.

Learning receptive fields using ℓ_1 sparsity

- ▶ To learn the basis functions $\{\vec{b}_i\}$ from a set of natural images $\{\vec{I}^\mu : \mu \in \Lambda\}$, we extend equation (6) to obtain a criteria $E(b, \alpha)$ for fitting basis functions b and coefficients α to the set of images:

$$E(b, \alpha) = \sum_{\mu \in \Lambda} (I^\mu(\vec{x}) - \sum_{i=1}^N \alpha_i^\mu b_i(\vec{x}))^2 + \lambda \sum_{\mu \in \Lambda} \sum_{i=1}^N |\alpha_i^\mu|.$$

- ▶ We estimate the basis functions \hat{b} and the coefficients $\hat{\alpha}$ by minimizing $E(b, \alpha)$ to obtain:

$$(\hat{b}, \hat{\alpha}) = \arg \min_{(b, \alpha)} E(b, \alpha).$$

- ▶ This criterion has been applied to natural images (where the \vec{I} represent small image regions), and the resulting basis functions, see figure (18)(left), include filters that look like Gabor functions but they also include other types of filters observed in experiments (Olshausen, 1996).

Alternatives: ICA

- ▶ Other methods can predict receptive field properties from natural images using a similar image model, $I(\vec{x}) = \sum_{i=1}^N \alpha_i b_i(\vec{x})$, but imposing different assumptions on the form of the bases. In particular, independent component analysis (ICA) gives similar receptive field models (van Hateren & Ruderman, 1998). Hyvarinen (2010) explains this by showing that both types of models – L1 sparsity and ICA – encourage the α_i to be strongly peaked at 0, but can occasionally have large nonzero values.
- ▶ What happens if we remove the sparsity requirement and instead find the basis functions that minimize $\sum_{\mu \in \Lambda} (I^\mu(\vec{x}) - \sum_{i=1}^N \alpha_i^\mu b_i(\vec{x}))^2$? The basis functions will be the eigenvectors of the correlation matrix of the images and can be found by principal component analysis (PCA).

ℓ_1 sparsity figure

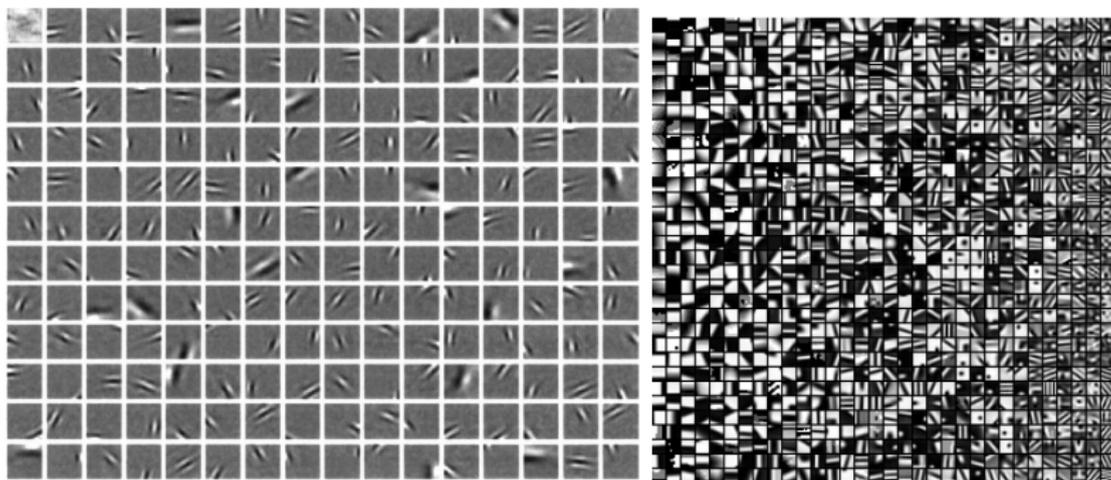


Figure 18 : Left: The receptive fields learned using sparsity (Olshausen, 1996).
Right: receptive fields learned by matched filters.

Matched filter interpretation

- ▶ An alternative idea is that cells are feature detectors (Lettvin et al., 1959). This can be modelled by a set of matched filters, which is an extreme form of sparsity, because any image patch can be represented by a single filter.
- ▶ Examples of matched filters are shown in the previous figure (right).
- ▶ Suppose we have a filter \vec{W} and an input image patch \vec{I}_p . We want to find the best fit of the filter to the image by allowing us to transform the filter by $\vec{W} \mapsto a\vec{W} + b\vec{e}$, where $\vec{e} = (1/\sqrt{N})(1, \dots, 1)$. This corresponds to scaling the filter by a and adding a constant vector b . If \vec{W} is a derivative filter, then by definition, $\vec{W} \cdot \vec{e} = 0$. We normalize \vec{W} and \vec{e} so that $\vec{W} \cdot \vec{W} = \vec{e} \cdot \vec{e} = 1$.

Matched filters

- ▶ The goal is to find the best scaling/contrast a and background b to minimize the match:

$$E(a, b) = |\vec{I}_p - a\vec{W} - b\vec{e}|^2.$$

- ▶ The solutions \hat{a}, \hat{b} are given by (take derivatives of E with respect to a and b , recalling that \vec{W} and \vec{e} are normalized):

$$\hat{a} = \vec{W} \cdot \vec{I}_p, \quad \hat{b} = \vec{e} \cdot \vec{I}_p.$$

- ▶ The filter response is just the best estimate of the contrast a . The estimate of the background b is just the mean value of the image. Finally, the energy $E(\hat{a}, \hat{b})$ is a measure of how well the filter “matches” the input image.

Matched filter dictionaries

The idea of a matched filter leads naturally to the idea of having a “dictionary” of filters $\{\vec{W}^\mu : \mu \in \Lambda\}$, in which different filters \vec{W}^μ are tuned to different types of image patches. In other words, the input image patch is encoded by the filter that best matches it. The dictionary of matched filters could be implemented by a set of cells (e.g., orientation columns). In this interpretation, the magnitude of the dot product $\vec{W} \cdot \vec{I}$ is less important than deciding which filter best matches the input \vec{I}_p . Matched filters can be thought of as an extreme case of sparsity. In the previous slides, an image was represented by a linear combination of basis functions whose weights were penalized by the ℓ_1 , $\sum_i |\alpha_i|$. By comparison, matched filters represent an image by a single basis function. This gives an ever sparser representation of the image, but at the possible cost of a much larger image dictionary. Matched filters can be thought of as *feature detectors* because they respond only to very specific inputs.