



HAL
open science

Sparse models for Computer Vision

Laurent Perrinet

► **To cite this version:**

Laurent Perrinet. Sparse models for Computer Vision. Biologically inspired computer vision, 2015, 9783527680863. 10.1002/9783527680863.ch14 . hal-01444362

HAL Id: hal-01444362

<https://hal.science/hal-01444362v1>

Submitted on 24 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - ShareAlike 4.0 International License

Chapter 14

Sparse models for Computer Vision

BibTex entry

This chapter appeared as (Perrinet, 2015):

```
@inbook{Perrinet15sparse,  
  Author = {Perrinet, Laurent U.},  
  Chapter = {Sparse models for Computer Vision},  
  Citeulike-Article-Id = {13514904},  
  Editor = {Crist\{'o\}bal, Gabriel and Keil, Matthias S. and Perrinet, Laurent U.},  
  Keywords = {bicv-sparse, sanz12jnp, vacher14},  
    month = nov,  
  chapter = {14},  
  Priority = {0},  
  isbn = {9783527680863},  
  title = {Sparse Models for Computer Vision},  
  DOI = {10.1002/9783527680863.ch14},  
    url={http://onlinelibrary.wiley.com/doi/10.1002/9783527680863.ch14/summary},  
    publisher = {Wiley-VCH Verlag GmbH {\&} Co. KGaA},  
  booktitle = {Biologically inspired computer vision},  
  Year = {2015}  
}
```


Contents

14 Sparse models for Computer Vision	1
14.1 Motivation	3
14.1.1 Efficiency and sparseness in biological representations of natural images	3
14.1.2 Sparseness induces neural organization	4
14.1.3 Outline: Sparse models for Computer Vision	5
14.2 What is sparseness? Application to image patches	6
14.2.1 Definitions of sparseness	6
14.2.2 Learning to be sparse: the SparseNet algorithm	8
14.2.3 Results: efficiency of different learning strategies	9
14.3 SparseLets: a multi scale, sparse, biologically inspired representation of natural images	10
14.3.1 Motivation: architecture of the primary visual cortex	10
14.3.2 The SparseLets framework	12
14.3.3 Efficiency of the SparseLets framework	14
14.4 SparseEdges: introducing prior information	17
14.4.1 Using the prior in first-order statistics of edges	17
14.4.2 Using the prior statistics of edge co-occurrences	19
14.5 Conclusion	22

14.1 Motivation

14.1.1 Efficiency and sparseness in biological representations of natural images

The central nervous system is a dynamical, adaptive organ which constantly evolves to provide optimal decisions¹ for interacting with the environment. The early visual pathways provides with a powerful system for probing and modeling these mechanisms. For instance, the primary visual cortex of primates (V1) is absolutely central for most visual tasks. There, it is observed that some neurons from the input layer of V1 present a selectivity for localized, edge-like features —as represented by their “receptive fields” (Hubel and Wiesel, 1968). Crucially, there is experimental evidence for *sparse* firing in the neocortex (Barth and Poulet, 2012; Willmore et al., 2011) and in particular in V1. A representation is sparse when each input signal is associated with a relatively small sub-set of simultaneously activated neurons within a whole population. For instance, orientation selectivity of simple cells is sharper than the selectivity that would be predicted by linear filtering. Such a procedure produces a rough “sketch” of the image on the surface of V1 that is believed to serve as a “blackboard” for higher-level cortical areas (Marr, 1983). However, it is still largely unknown how neural computations act in V1 to represent the image. More specifically, what is the role of sparseness —as a generic neural signature— in the global function of neural computations?

¹Decisions are defined in their broader sense of elementary choices operated in the system at associative or motor levels (Friston et al., 2012).

A popular view is that such a population of neurons operates such that relevant sensory information from the retino-thalamic pathway is transformed (or “coded”) efficiently. Such efficient representation will allow decisions to be taken optimally in higher-level areas. In this framework, optimality is defined in terms of information theory (Attneave, 1954; Atick, 1992; Wolfe et al., 2010). For instance, the representation produced by the neural activity in V1 is sparse: It is believed that this reduces redundancies and allows to better segregate edges in the image (Field, 1994; Froudarakis et al., 2014). This optimization is operated given biological constraints, such as the limited bandwidth of information transfer to higher processing stages or the limited amount of metabolic resources (energy or wiring length). More generally it allows to increase the storage capacity of associative memories before memory patterns start to interfere with each other (Palm, 2013). Moreover, it is now widely accepted that this redundancy reduction is achieved in a neural population through lateral interactions. Indeed, a link between anatomical data and a functional connectivity between neighboring representations of edges has been found (Bosking et al., 1997), though their conclusions were more recently refined to show that this process may be more complex (Hunt et al., 2011). By linking neighboring neurons representing similar features, one allows thus a more efficient representation in V1. As computer vision systems are subject to similar constraints, applying such a paradigm therefore seems a promising approach towards more biomimetic algorithms.

It is believed that such a property reflects the efficient match of the representation with the statistics of natural scenes, that is, to behaviorally relevant sensory inputs. Indeed, sparse representations are prominently observed for cortical responses to natural stimuli (Field, 1987; Vinje and Gallant, 2000; DeWeese et al., 2003; Baudot et al., 2013). As the function of neural systems mostly emerges from unsupervised learning, it follows that these are adapted to the input which are behaviorally the most common and important. More generally, by being adapted to natural scenes, this shows that sparseness is a neural signature of an underlying optimization process. In fact, one goal of neural computation in low-level sensory areas such as V1 is to provide relevant predictions (Rao and Ballard, 1999; Spratling, 2011). This is crucial for living beings as they are often confronted with noise (internal to the brain or external, such as in low light conditions), ambiguities (such as inferring a three dimensional slant from a bi-dimensional retinal image). Also, the system has to compensate for inevitable delays, such as the delay from light stimulation to activation in V1 which is estimated to be of 50 ms in humans. For instance, a tennis ball moving at 20 m s^{-1} at one meter in the frontal plane elicits an input activation in V1 corresponding to around 45° of visual angle behind its physical position (Perrinet et al., 2014). Thus, to be able to translate such knowledge to the computer vision community, it is crucial to better understand *why* the neural processes that produce sparse coding are efficient.

14.1.2 Sparseness induces neural organization

A breakthrough in the modeling of the representation in V1 was the discovery that sparseness is sufficient to induce the emergence of receptive fields similar to V1 simple cells (Olshausen and Field, 1996). This reflects the fact that, at the learning time scale, coding is optimized relative to the statistics of natural scenes such that independent components of the input are represented (Olshausen and Field, 1997; Bell and Sejnowski, 1997). The emergence of edge-like simple cell receptive fields in the input layer of area V1 of primates may thus be considered as a coupled coding and learning optimization problem: At the coding time scale, the sparseness of the representation is optimized for any given input while at the learning time scale, synaptic weights are tuned to achieve on average an optimal representation efficiency over natural scenes. This theory has allowed to connect the different fields by providing a link between information theory models, neuromimetic models and physiological observations.

In practice, most sparse unsupervised learning models aim at optimizing a cost defined on prior assumptions on the sparseness of the representation. These sparse learning algorithms have been

applied both for images (Fyfe and Baddeley, 1995; Olshausen and Field, 1996; Zibulevsky and Pearlmuter, 2001; Perrinet et al., 2004; Rehn and Sommer, 2007; Doi et al., 2007; Perrinet, 2010) and sounds (Lewicki and Sejnowski, 2000; Smith and Lewicki, 2006). Sparse coding may also be relevant to the amount of energy the brain needs to use to sustain its function. The total neural activity generated in a brain area is inversely related to the sparseness of the code, therefore the total energy consumption decreases with increasing sparseness. As a matter of fact the probability distribution functions of neural activity observed experimentally can be approximated by so-called exponential distributions, which have the property of maximizing information transmission for a given mean level of activity (Baddeley et al., 1997). To solve such constraints, some models thus directly compute a sparseness cost based on the representation’s distribution. For instance, the kurtosis corresponds to the 4th statistical moment (the first three moments being in order the mean, variance and skewness) and measures how the statistics deviates from a Gaussian: A positive kurtosis measures if this distribution has an “heavier tail” than a Gaussian for a similar variance — and thus corresponds to a sparser distribution. Based on such observations, other similar statistical measures of sparseness have been derived in the neuroscience literature (Vinje and Gallant, 2000).

A more general approach is to derive a representation cost. For instance, learning is accomplished in the SPARSENET algorithmic framework (Olshausen and Field, 1997) on image patches taken from natural images as a sequence of coding and learning steps. First, sparse coding is achieved using a gradient descent over a convex cost. We will see later in this chapter how this cost is derived from a prior on the probability distribution function of the coefficients and how it favors the sparseness of the representation. At this step, the coding is performed using the current state of the “dictionary” of receptive fields. Then, knowing this sparse solution, learning is defined as slowly changing the dictionary using Hebbian learning (Hebb, 1949). As we will see later, the parameterization of the prior has a major impact on the results of the sparse coding and thus on the emergence of edge-like receptive fields and requires proper tuning. Yet, this class of models provides a simple solution to the problem of sparse representation in V1.

However, these models are quite abstract and assume that neural computations may estimate some rather complex measures such as gradients - a problem that may also be faced by neuromorphic systems. Efficient, realistic implementations have been proposed which show that imposing sparseness may indeed guide neural organization in neural network models, see for instance (Zylberberg et al., 2011; Hunt et al., 2013). Additionally, it has also been shown that in a neuromorphic model, an efficient coding hypothesis links sparsity and selectivity of neural responses (Blättler and Hahnloser, 2011). More generally, such neural signatures are reminiscent of the shaping of neural activity to account for contextual influences. For instance, it is observed that —depending on the context outside the receptive field of a neuron in area V1— the tuning curve may demonstrate a modulation of its orientation selectivity. This was accounted for instance as a way to optimize the coding efficiency of a population of neighboring neurons (Serìès et al., 2004). As such, sparseness is a relevant neural signature for a large class of neural computations implementing efficient coding.

14.1.3 Outline: Sparse models for Computer Vision

As a consequence, sparse models provide a fruitful approach for computer vision. It should be noted that other popular approaches for taking advantage of sparse representations exist. The most popular is compressed sensing (Ganguli and Sompolinsky, 2012), for which it has been proven that —assuming sparseness in the input, it is possible to reconstruct the input from a sparse choice of linear coefficients computed from randomly drawn basis functions. Note also that some studies also focus on temporal sparseness. Indeed, by computing for a given neuron the relative numbers of active events relative to a given time window, one computes the so-called lifetime sparseness (see for instance (Willmore et al., 2011)). We will see below that this measure may be related to population sparseness. For a review of sparse modeling approaches, we refer to (Elad, 2010). Herein, we will

focus on the particular sub-set of such models based on their biological relevance.

Indeed, we will rather focus on biomimetic sparse models as tools to shape future computer vision algorithms (Benoit et al., 2010; Serre and Poggio, 2010). In particular, we will not review models which mimic neural activity, but rather on algorithms which mimic their efficiency, bearing in mind the constraints that are linked to neural systems (no central clock, internal noise, parallel processing, metabolic cost, wiring length). For that purpose, we will complement some previous studies (Perrinet et al., 2004; Fischer et al., 2007a; Perrinet, 2008, 2010) (for a review see (Perrinet and Masson, 2007)) by putting these results in light of most recent theoretical and physiological findings.

This chapter is organized as follows. First, in Section 14.2 we will outline how we may implement the unsupervised learning algorithm at a local scale for image patches. Then we will extend in Section 14.3 such an approach to full scale natural images by defining the *SparseLets* framework. Such formalism will then be extended in Section 14.4 to include context modulation, for instance from higher-order areas. These different algorithms (from the local scale of image patches to more global scales) will each be accompanied by a supporting implementation (with the source code) for which we will show example usage and results. We will in particular highlight novel results and then draw some conclusions on the perspective of sparse models for computer vision. More specifically, we will propose that bio-inspired approaches may be applied to computer vision using predictive coding schemes, sparse models being one simple and efficient instance of such schemes.

14.2 What is sparseness? Application to image patches

14.2.1 Definitions of sparseness

In low-level sensory areas, the goal of neural computations is to generate efficient intermediate *representations* as we have seen that this allows more efficient decision making. Classically, a representation is defined as the inversion of an internal generative model of the sensory world, that is, by inferring the sources that generated the input signal. Formally, as in (Olshausen and Field, 1997), we define a Generative Linear Model (GLM) for describing natural, static, grayscale image patches \mathbf{I} (represented by column vectors of dimension L pixels), by setting a “dictionary” of M images (also called “atoms” or “filters”) as the $L \times M$ matrix $\Phi = \{\Phi_i\}_{1 \leq i \leq M}$. Knowing the associated “sources” as a vector of coefficients $\mathbf{a} = \{a_i\}_{1 \leq i \leq M}$, the image is defined using matrix notation as a sum of weighted atoms:

$$\mathbf{I} = \Phi \mathbf{a} + \mathbf{n} \quad (14.1)$$

where \mathbf{n} is a Gaussian additive noise image. This noise, as in (Olshausen and Field, 1997), is scaled to a variance of σ_n^2 to achieve decorrelation by applying Principal Component Analysis to the raw input images, without loss of generality since this preprocessing is invertible. Generally, the dictionary Φ may be much larger than the dimension of the input space (that is, if $M \gg L$) and it is then said to be *over-complete*. However, given an over-complete dictionary, the inversion of the GLM leads to a combinatorial search and typically, there may exist many coding solutions \mathbf{a} to Eq. 14.1 for one given input \mathbf{I} . The goal of efficient coding is to find, given the dictionary Φ and for any observed signal \mathbf{I} , the “best” representation vector, that is, as close as possible to the sources that generated the signal. Assuming that for simplicity, each individual coefficient is represented in the neural activity of a single neuron, this would justify the fact that this activity is sparse. It is therefore necessary to define an efficiency criterion in order to choose between these different solutions.

Using the GLM, we will infer the “best” coding vector as the most probable. In particular, from the physics of the synthesis of natural images, we know *a priori* that image representations are sparse: They are most likely generated by a small number of features relatively to the dimension M of the representation space. Similarly to Lewicki and Sejnowski (2000), this can be formalized in the probabilistic framework defined by the GLM (see Eq. 14.1), by assuming that knowing the

prior distribution of the coefficients a_i for natural images, the representation cost of \mathbf{a} for one given natural image is:

$$\begin{aligned} \mathcal{C}(\mathbf{a}|\mathbf{I}, \Phi) &\stackrel{\text{def}}{=} -\log P(\mathbf{a}|\mathbf{I}, \Phi) = \log P(\mathbf{I}) - \log P(\mathbf{I}|\mathbf{a}, \Phi) - \log P(\mathbf{a}|\Phi) \\ &= \log P(\mathbf{I}) + \frac{1}{2\sigma_n^2} \|\mathbf{I} - \Phi\mathbf{a}\|^2 - \sum_i \log P(a_i|\Phi) \end{aligned} \quad (14.2)$$

where $P(\mathbf{I})$ is the partition function which is independent of the coding (and that we thus ignore in the following) and $\|\cdot\|$ is the L_2 -norm in image space. This efficiency cost is measured in bits if the logarithm is of base 2, as we will assume without loss of generality thereafter. For any representation \mathbf{a} , the cost value corresponds to the description length (Rissanen, 1978): On the right hand side of Eq. 14.2, the second term corresponds to the information from the image which is not coded by the representation (reconstruction cost) and thus to the information that can be at best encoded using entropic coding pixel by pixel (that is, the negative log-likelihood $-\log P(\mathbf{I}|\mathbf{a}, \Phi)$ in Bayesian terminology, see chapter 009_series for Bayesian models applied to computer vision). The third term $S(\mathbf{a}|\Phi) = -\sum_i \log P(a_i|\Phi)$ is the representation or sparseness cost: It quantifies representation efficiency as the coding length of each coefficient of \mathbf{a} which would be achieved by entropic coding knowing the prior and assuming that they are independent. The rightmost penalty term (see Equation 14.2) gives thus a definition of sparseness $S(\mathbf{a}|\Phi)$ as the sum of the log prior of coefficients.

In practice, the sparseness of coefficients for natural images is often defined by an *ad hoc* parameterization of the shape of the prior. For instance, the parameterization in Olshausen and Field (1997) yields the coding cost:

$$\mathcal{C}_1(\mathbf{a}|\mathbf{I}, \Phi) = \frac{1}{2\sigma_n^2} \|\mathbf{I} - \Phi\mathbf{a}\|^2 + \beta \sum_i \log\left(1 + \frac{a_i^2}{\sigma^2}\right) \quad (14.3)$$

where β corresponds to the steepness of the prior and σ to its scaling (see Figure 13.2 from (Olshausen, 2002)). This choice is often favored because it results in a convex cost for which known numerical optimization methods such as conjugate gradient may be used. In particular, these terms may be put in parallel to regularization terms that are used in computer vision. For instance, a L_2 -norm penalty term corresponds to Tikhonov regularization (Tikhonov, 1977) or a L_1 -norm term corresponds to the Lasso method. See chapter 003_holly_gerhard for a review of possible parametrization of this norm, for instance by using nested L_p norms. Classical implementation of sparse coding rely therefore on a parametric measure of sparseness.

Let's now derive another measure of sparseness. Indeed, a non-parametric form of sparseness cost may be defined by considering that neurons representing the vector \mathbf{a} are either active or inactive. In fact, the spiking nature of neural information demonstrates that the transition from an inactive to an active state is far more significant at the coding time scale than smooth changes of the firing rate. This is for instance perfectly illustrated by the binary nature of the neural code in the auditory cortex of rats (DeWeese et al., 2003). Binary codes also emerge as optimal neural codes for rapid signal transmission (Bethge et al., 2003; Nikitin et al., 2009). This is also relevant for neuromorphic systems which transmit discrete events (such as a network packet). With a binary event-based code, the cost is only incremented when a new neuron gets active, regardless to the analog value. Stating that an active neuron carries a bounded amount of information of λ bits, an upper bound for the representation cost of neural activity on the receiver end is proportional to the count of active neurons, that is, to the ℓ_0 pseudo-norm $\|\mathbf{a}\|_0$:

$$\mathcal{C}_0(\mathbf{a}|\mathbf{I}, \Phi) = \frac{1}{2\sigma_n^2} \|\mathbf{I} - \Phi\mathbf{a}\|^2 + \lambda \|\mathbf{a}\|_0 \quad (14.4)$$

This cost is similar with information criteria such as the Akaike Information Criteria (Akaike, 1974) or distortion rate (Mallat, 1998, p. 571). This simple non-parametric cost has the advantage of being

dynamic: The number of active cells for one given signal grows in time with the number of spikes reaching the target population. But Eq. 14.4 defines a harder cost to optimize (in comparison to Equation 14.3 for instance) since the hard ℓ_0 pseudo-norm sparseness leads to a non-convex optimization problem which is *NP-complete* with respect to the dimension M of the dictionary (Mallat, 1998, p. 418).

14.2.2 Learning to be sparse: the SparseNet algorithm

We have seen above that we may define different models for measuring sparseness depending on our prior assumption on the distribution of coefficients. Note first that, assuming that the statistics are stationary (more generally ergodic), then these measures of sparseness across a population should necessarily imply a lifetime sparseness for any neuron. Such a property is essential to extend results from electro-physiology. Indeed, it is easier to record a restricted number of cells than a full population (see for instance (Willmore et al., 2011)). However, the main property in terms of efficiency is that the representation should be sparse at any given time, that is, in our setting, at the presentation of each novel image.

Now that we have defined sparseness, how could we use it to induce neural organization? Indeed, given a sparse coding strategy that optimizes any representation efficiency cost as defined above, we may derive an unsupervised learning model by optimizing the dictionary Φ over natural scenes. On the one hand, the flexibility in the definition of the sparseness cost leads to a wide variety of proposed *sparse coding* solutions (for a review, see (Pece, 2002)) such as numerical optimization (Olshausen and Field, 1997), non-negative matrix factorization (Lee and Seung, 1999; Ranzato et al., 2007) or Matching Pursuit (Perrinet et al., 2004; Smith and Lewicki, 2006; Rehn and Sommer, 2007; Perrinet, 2010). They are all derived from correlation-based inhibition since this is necessary to remove redundancies from the linear representation. This is consistent with the observation that lateral interactions are necessary for the formation of elongated receptive fields (Bolz and Gilbert, 1989; Wolfe et al., 2010).

On the other hand, these methods share the same GLM model (see Eq. 14.1) and once the sparse coding algorithm is chosen, the learning scheme is similar. As a consequence, after every coding sweep, we increased the efficiency of the dictionary Φ with respect to Eq. 14.2. This is achieved using the online gradient descent approach given the current sparse solution, $\forall i$:

$$\Phi_i \leftarrow \Phi_i + \eta \cdot a_i \cdot (\mathbf{I} - \Phi \mathbf{a}) \quad (14.5)$$

where η is the learning rate. Similarly to Eq. 17 in (Olshausen and Field, 1997) or to Eq. 2 in (Smith and Lewicki, 2006), the relation is a linear “Hebbian” rule (Hebb, 1949) since it enhances the weight of neurons proportionally to the correlation between pre- and post-synaptic neurons. Note that there is no learning for non-activated coefficients. The novelty of this formulation compared to other linear Hebbian learning rule such as (Oja, 1982) is to take advantage of the sparse representation, hence the name Sparse Hebbian Learning (SHL).

The class of SHL algorithms are unstable without homeostasis, that is, without a process that maintains the system in a certain equilibrium. In fact, starting with a random dictionary, the first filters to learn are more likely to correspond to salient features (Perrinet et al., 2004) and are therefore more likely to be selected again in subsequent learning steps. In SPARSENET, the homeostatic gain control is implemented by adaptively tuning the norm of the filters. This method equalizes the variance of coefficients across neurons using a geometric stochastic learning rule. The underlying heuristic is that this introduces a bias in the choice of the active coefficients. In fact, if a neuron is not selected often, the geometric homeostasis will decrease the norm of the corresponding filter, and therefore —from Eq. 14.1 and the conjugate gradient optimization— this will increase the value of the associated scalar. Finally, since the prior functions defined in Eq. 14.3 are identical for all neurons, this will increase the relative probability that the neuron is selected with a higher relative

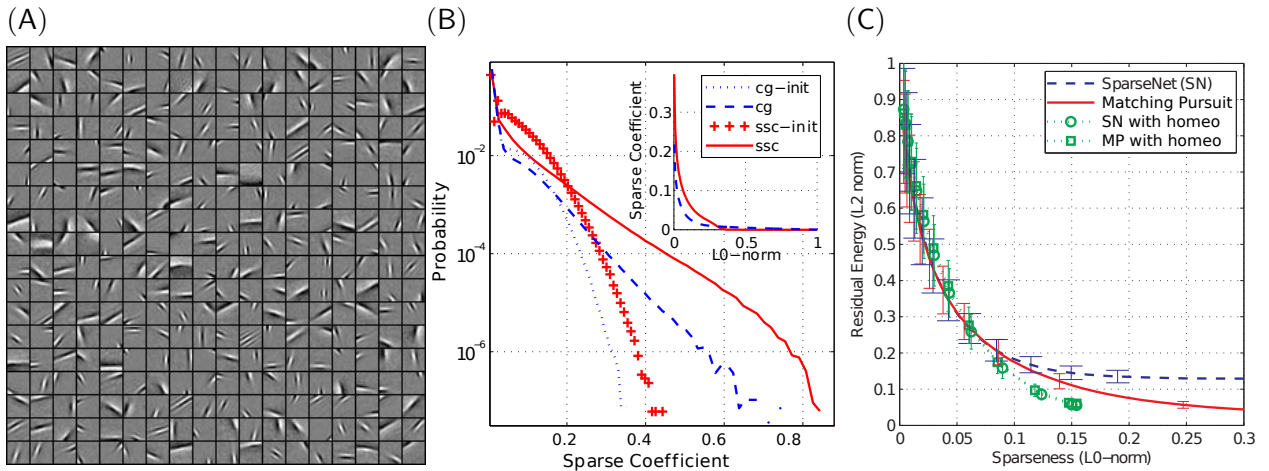


Figure 14.1: **Learning a sparse code using Sparse Hebbian Learning.** (A) We show the results at convergence (20000 learning steps) of a sparse model with unsupervised learning algorithm which progressively optimizes the relative number of active (non-zero) coefficients (ℓ_0 pseudo-norm) (Perrinet, 2010). Filters of the same size as the image patches are presented in a matrix (separated by a black border). Note that their position in the matrix is arbitrary as in ICA. These results show that sparseness induces the emergence of edge-like receptive fields similar to those observed in the primary visual area of primates. (B) We show the probability distribution function of sparse coefficients obtained by our method compared to (Olshausen and Field, 1996) with first, random dictionaries (respectively 'ssc-init' and 'cg-init') and second, with the dictionaries obtained after convergence of respective learning schemes (respectively 'ssc' and 'cg'). At convergence, sparse coefficients are more sparsely distributed than initially, with more kurtotic probability distribution functions for 'ssc' in both cases, as can be seen in the “longer tails” of the distribution. (C) We evaluate the coding efficiency of both methods by plotting the average residual error (L_2 norm) as a function of the ℓ_0 pseudo-norm. This provides a measure of the coding efficiency for each dictionary over the set of image patches (error bars represent one standard deviation). Best results are those providing a lower error for a given sparsity (better compression) or a lower sparseness for the same error.

value. The parameters of this homeostatic rule have a great importance for the convergence of the global algorithm. In (Perrinet, 2010), we have derived a more general homeostasis mechanism derived from the optimization of the representation efficiency through histogram equalization which we will describe later (see Section 14.4.1).

14.2.3 Results: efficiency of different learning strategies

The different SHL algorithms simply differ by the coding step. This implies that they only differ by first, how sparseness is defined at a functional level and second, how the inverse problem corresponding to the coding step is solved at the algorithmic level. Most of the schemes cited above use a less strict, parametric definition of sparseness (like the convex L_1 -norm), but for which a mathematical formulation of the optimization problem exists. Few studies such as (Liu and Jia, 2014; Peharz and Pernkopf, 2012) use the stricter ℓ_0 pseudo-norm as the coding problem gets more difficult. A thorough comparison of these different strategies was recently presented in (Charles et al., 2012). See also (Aharon et al., 2006) for properties of the coding solutions to the ℓ_0 pseudo-norm. Similarly, in (Perrinet, 2010), we preferred to retrieve an approximate solution to the coding problem to have a better match with the measure of efficiency Eq. 14.4.

Such an algorithmic framework is implemented in the SHL-scripts package². These scripts allow the retrieval of the database of natural images and the replication of the results of (Perrinet, 2010) reported in this section. With a correct tuning of parameters, we observed that different coding schemes show qualitatively a similar emergence of edge-like filters. The specific coding algorithm used to obtain this sparseness appears to be of secondary importance as long as it is adapted to the data and yields sufficiently efficient sparse representation vectors. However, resulting dictionaries vary qualitatively among these schemes and it was unclear which algorithm is the most efficient and what was the individual role of the different mechanisms that constitute SHL schemes. At the learning level, we have shown that the homeostasis mechanism had a great influence on the qualitative distribution of learned filters (Perrinet, 2010).

Results are shown in Figure 14.1. This figure represents the qualitative results of the formation of edge-like filters (receptive fields). More importantly, it shows the quantitative results as the average decrease of the squared error as a function of the sparseness. This gives a direct access to the cost as computed in Equation 14.4. These results are comparable with the SPARSENET algorithm. Moreover, this solution, by giving direct access to the atoms (filters) that are chosen, provides with a more direct tool to manipulate sparse components. One further advantage consists in the fact that this unsupervised learning model is non-parametric (compare with Eq. 14.3) and thus does not need to be parametrically tuned. Results show the role of homeostasis on the unsupervised algorithm. In particular, using the comparison of coding and decoding efficiency with and without this specific homeostasis, we have proven that cooperative homeostasis optimized overall representation efficiency (see also Section 14.4.1).

It is at this point important to note that in this algorithm, we achieve an exponential convergence of the squared error (Mallat, 1998, p. 422), but also that this curve can be directly derived from the coefficients' values. Indeed, for N coefficients (that is, $\|\mathbf{a}\|_0 = N$), we have the squared error equal to:

$$E_N \stackrel{\text{def}}{=} \|\mathbf{I} - \Phi\mathbf{a}\|^2 / \|\mathbf{I}\|^2 = 1 - \sum_{1 \leq k \leq N} a_k^2 / \|\mathbf{I}\|^2 \quad (14.6)$$

As a consequence, the sparser the distributions of coefficients, then quicker is the decrease of the residual energy. In the following section, we will describe different variations of this algorithm. To compare their respective efficiency, we will plot the decrease of the coefficients along with the decrease of the residual's energy. Using such tools, we will now explore if such a property extends to full-scale images and not only to image patches, an important condition for using sparse models in computer vision algorithms.

14.3 SparseLets: a multi scale, sparse, biologically inspired representation of natural images

14.3.1 Motivation: architecture of the primary visual cortex

Our goal here is to build practical algorithms of sparse coding for computer vision. We have seen above that it is possible to build an adaptive model of sparse coding that we applied to 12×12 image patches. Invariably, this has shown that the independent components of image patches are edge-like filters, such as is found in simple cells of V1. This model has shown that for randomly chosen image patches, these may be described by a sparse vector of coefficients. Extending this result to full-field natural images, we can expect that this sparseness would increase by a degree of order. In fact, except in a densely cluttered image such as a close-up of a texture, natural images tend to have wide areas which are void (such as the sky, walls or uniformly filled areas). However, applying directly

²These scripts are available at https://github.com/bicv/SHL_scripts and documented at https://pythonhosted.org/SHL_scripts.

the SPARSENET algorithm to full-field images is impossible in practice as its computer simulation would require too much memory to store the over-complete set of filters. However, it is still possible to define *a priori* these filters and herein, we will focus on a full-field sparse coding method whose filters are inspired by the architecture of the primary visual cortex.

The first step of our method involves defining the dictionary of templates (or filters) for detecting edges. We use a log-Gabor representation, which is well suited to represent a wide range of natural images (Fischer et al., 2007a). This representation gives a generic model of edges parameterized by their shape, orientation, and scale. We set the range of these parameters to match with what has been reported for simple-cell responses in macaque primary visual cortex (V1). Indeed log-Gabor filters are similar to standard Gabors and both are well fitted to V1 simple cells (Daugman, 1980). Log-Gabors are known to produce a sparse set of linear coefficients (Field, 1999). Like Gabors, these filters are defined by Gaussians in Fourier space, but their specificity is that log-Gabors have Gaussians envelopes in log-polar frequency space. This is consistent with physiological measurements which indicate that V1 cell responses are symmetric on the log frequency scale. They have multiple advantages over Gaussians: In particular, they have no DC component, and more generally, their envelopes more broadly cover the frequency space (Fischer et al., 2007b). In this chapter, we set the bandwidth of the Fourier representation of the filters to 0.4 and $\pi/8$ respectively in the log-frequency and polar coordinates to get a family of relatively elongated (and thus selective) filters (see Fischer et al. (2007b) and Figure 14.2-A for examples of such edges). Prior to the analysis of each image, we used the spectral whitening filter described by Olshausen and Field (1997) to provide a good balance of the energy of output coefficients (Perrinet et al., 2004; Fischer et al., 2007a). Such a representation is implemented in the `LogGabor` package³.

This transform is linear and can be performed by a simple convolution repeated for every edge type. Following Fischer et al. (2007b), convolutions were performed in the Fourier (frequency) domain for computational efficiency. The Fourier transform allows for a convenient definition of the edge filter characteristics, and convolution in the spatial domain is equivalent to a simple multiplication in the frequency domain. By multiplying the envelope of the filter and the Fourier transform of the image, one may obtain a filtered spectral image that may be converted to a filtered spatial image using the inverse Fourier transform. We exploited the fact that by omitting the symmetrical lobe of the envelope of the filter in the frequency domain, we obtain quadrature filters. Indeed, the output of this procedure gives a complex number whose real part corresponds to the response to the symmetrical part of the edge, while the imaginary part corresponds to the asymmetrical part of the edge (see Fischer et al. (2007b) for more details). More generally, the modulus of this complex number gives the energy response to the edge —as can be compared to the response of complex cells in area V1, while its argument gives the exact phase of the filter (from symmetric to non-symmetric). This property further expands the richness of the representation.

Given a filter at a given orientation and scale, a linear convolution model provides a translation-invariant representation. Such invariance can be extended to rotations and scalings by choosing to multiplex these sets of filters at different orientations and spatial scales. Ideally, the parameters of edges would vary in a continuous fashion, to a full relative translation, rotation, and scale invariance. However this is difficult to achieve in practice and some compromise has to be found. Indeed, though orthogonal representations are popular in computer vision due to their computational tractability, it is desirable in our context that we have a relatively high over-completeness in the representation to achieve this invariance. For a given set of 256×256 images, we first chose to have 8 dyadic levels (that is, by doubling the scale at each level) with 24 different orientations. Orientations are measured as a non-oriented angle in radians, by convention in the range from $-\frac{\pi}{2}$ to $\frac{\pi}{2}$ (but not including $-\frac{\pi}{2}$) with respect to the x -axis. Finally, each image is transformed into a pyramid of coefficients. This

³Python scripts are available at <https://github.com/bicv/LogGabor> and documented at <https://pythonhosted.org/LogGabor>.

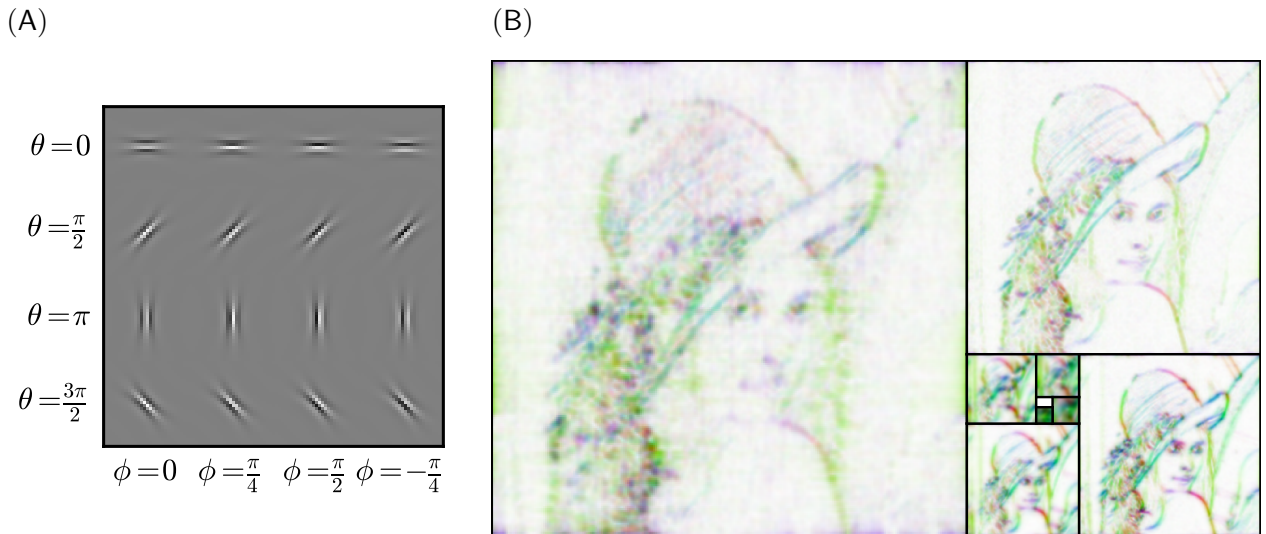


Figure 14.2: **The log-Gabor pyramid** (A) A set of log-Gabor filters showing first in different rows different orientations and second in different columns different phases. Here we have only shown one scale. Note the similarity with Gabor filters. (B) Using this set of filter, one can define a linear representation that is rotation-, scaling- and translation- invariant. Here we show a tiling of the different scales according to a Golden Pyramid (Perrinet, 2008). The hue gives the orientation while the value gives the absolute value (white denotes a low coefficient). Note the redundancy of the linear representation, for instance at different scales.

pyramid consists of approximately $4/3 \times 256^2 \approx 8.7 \times 10^4$ pixels multiplexed on 8 scales and 24 orientations, that is, approximately 16.7×10^6 coefficients, an over-completeness factor of about 256. This linear transformation is represented by a pyramid of coefficient, see Figure 14.2-B.

14.3.2 The SparseLets framework

The resulting dictionary of edge filters is over-complete. The linear representation would thus give a dense, relatively inefficient representation of the distribution of edges, see Figure 14.2-B. Therefore, starting from this linear representation, we searched instead for the most sparse representation. As we saw above in Section 14.2, minimizing the ℓ_0 pseudo-norm (the number of non-zero coefficients) leads to an expensive combinatorial search with regard to the dimension of the dictionary (it is NP-hard). As proposed by Perrinet et al. (2004), we may approximate a solution to this problem using a greedy approach. Such an approach is based on the physiology of V1. Indeed, it has been shown that inhibitory interneurons decorrelate excitatory cells to drive sparse code formation (Bolz and Gilbert, 1989; King et al., 2013). We use this local architecture to iteratively modify the linear representation (Fischer et al., 2007a).

In general, a greedy approach is applied when the optimal combination is difficult to solve globally, but can be solved progressively, one element at a time. Applied to our problem, the greedy approach corresponds to first choosing the single filter Φ_i that best fits the image along with a suitable coefficient a_i , such that the single source $a_i\Phi_i$ is a good match to the image. Examining every filter Φ_j , we find the filter Φ_i with the maximal correlation coefficient (“Matching” step), where:

$$i = \operatorname{argmax}_j \left(\left\langle \frac{\mathbf{I}}{\|\mathbf{I}\|}, \frac{\Phi_j}{\|\Phi_j\|} \right\rangle \right), \quad (14.7)$$

$\langle \cdot, \cdot \rangle$ represents the inner product, and $\|\cdot\|$ represents the L_2 (Euclidean) norm. The index (“address”) i gives the position (x and y), scale and orientation of the edge. We saw above that since filters at a given scale and orientation are generated by a translation, this operation can be efficiently computed using a convolution, but we keep this notation for its generality. The associated coefficient is the scalar projection:

$$a_i = \left\langle \mathbf{I}, \frac{\Phi_i}{\|\Phi_i\|^2} \right\rangle \quad (14.8)$$

Second, knowing this choice, the image can be decomposed as

$$\mathbf{I} = a_i \Phi_i + \mathbf{R} \quad (14.9)$$

where \mathbf{R} is the residual image (“Pursuit” step). We then repeat this 2-step process on the residual (that is, with $\mathbf{I} \leftarrow \mathbf{R}$) until some stopping criterion is met. Note also that the norm of the filters has no influence in this algorithm on the matching step or on the reconstruction error. For simplicity and without loss of generality, we will thereafter set the norm of the filters to 1: $\forall j, \|\Phi_j\| = 1$ (that is, that the spectral energy sums to 1). Globally, this procedure gives us a sequential algorithm for reconstructing the signal using the list of sources (filters with coefficients), which greedily optimizes the ℓ_0 pseudo-norm (i.e., achieves a relatively sparse representation given the stopping criterion). The procedure is known as the Matching Pursuit (MP) algorithm (Mallat and Zhang, 1993), which has been shown to generate good approximations for natural images (Perrinet et al., 2004; Perrinet, 2010).

We have included two minor improvements over this method: First, we took advantage of the response of the filters as complex numbers. As stated above, the modulus gives a response independent of the phase of the filter, and this value was used to estimate the best match of the residual image with the possible dictionary of filters (Matching step). Then, the phase was extracted as the argument of the corresponding coefficient and used to feed back onto the image in the Pursuit step. This modification allows for a phase-independent detection of edges, and therefore for a richer set of configurations, while preserving the precision of the representation.

Second, we used a “smooth” Pursuit step. In the original form of the Matching Pursuit algorithm, the projection of the Matching coefficient is fully removed from the image, which allows for the optimal decrease of the energy of the residual and allows for the quickest convergence of the algorithm with respect to the ℓ_0 pseudo-norm (i.e., it rapidly achieves a sparse reconstruction with low error). However, this efficiency comes at a cost, because the algorithm may result in non-optimal representations due to choosing edges sequentially and not globally. This is often a problem when edges are aligned (e.g. on a smooth contour), as the different parts will be removed independently, potentially leading to a residual with gaps in the line. Our goal here is not necessarily to get the fastest decrease of energy, but rather to provide with the best representation of edges along contours. We therefore used a more conservative approach, removing only a fraction (denoted by α) of the energy at each pursuit step (for MP, $\alpha = 1$). Note that in that case, Equation 14.6 has to be modified to account for the α parameter:

$$E_N = 1 - \alpha \cdot (2 - \alpha) \cdot \sum_{1 \leq k \leq N} a_k^2 / \|\mathbf{I}\|^2 \quad (14.10)$$

We found that $\alpha = 0.8$ was a good compromise between rapidity and smoothness. One consequence of using $\alpha < 1$ is that, when removing energy along contours, edges can overlap; even so, the correlation is invariably reduced. Higher and smaller values of α were also tested, and gave representation results similar to those presented here.

In summary, the whole coding algorithm is given by the following nested loops in pseudo-code:

1. draw a signal \mathbf{I} from the database; its energy is $E = \|\mathbf{I}\|^2$,

2. initialize sparse vector \mathbf{s} to zero and linear coefficients $\forall j, a_j = \langle \mathbf{I}, \Phi_j \rangle$,
3. while the residual energy $E = \|\mathbf{I}\|^2$ is above a given threshold do:
 - (a) select the best match: $i = \text{ArgMax}_j |a_j|$, where $|\cdot|$ denotes the modulus,
 - (b) increment the sparse coefficient: $s_i = s_i + \alpha \cdot a_i$,
 - (c) update residual image: $\mathbf{I} \leftarrow \mathbf{I} - \alpha \cdot a_i \cdot \Phi_i$,
 - (d) update residual coefficients: $\forall j, a_j \leftarrow a_j - \alpha \cdot a_i \langle \Phi_i, \Phi_j \rangle$,
4. the final set of non-zero values of the sparse representation vector \mathbf{s} constitutes the list of edges representing the image as the list of couples $\pi_i = (i, s_i)$, where i represents an edge occurrence as represented by its position, orientation and scale and s_i the complex-valued sparse coefficient.

This class of algorithms gives a generic and efficient representation of edges, as illustrated by the example in Figure 14.3-A. We also verified that the dictionary used here is better adapted to the extraction of edges than Gabors (Fischer et al., 2007a). The performance of the algorithm can be measured quantitatively by reconstructing the image from the list of extracted edges. All simulations were performed using Python (version 2.7.8) with packages NumPy (version 1.8.1) and SciPy (version 0.14.0) (Oliphant, 2007) on a cluster of Linux computing nodes. Visualization was performed using Matplotlib (version 1.3.1) (Hunter, 2007).⁴

14.3.3 Efficiency of the SparseLets framework

Figure 14.3-A, shows the list of edges extracted on a sample image. It fits qualitatively well with a rough sketch of the image. To evaluate the algorithm quantitatively, we measured the ratio of extracted energy in the images as a function of the number of edges on a database of 600 natural images of size 256×256^5 , see Figure 14.3-B. Measuring the ratio of extracted energy in the images, $N = 2048$ edges were enough to extract an average of approximately 97% of the energy of images in the database. To compare the strength of the sparseness in these full-scale images compared to the image patches discussed above (see Section 14.2), we measured the sparseness values obtained in images of different sizes. To be comparable, we measured the efficiency with respect to the relative ℓ_0 pseudo-norm in bits per unit of image surface (pixel): This is defined as the ratio of active coefficients times the numbers of bits required to code for each coefficient (that is, $\log_2(M)$, where M is total number of coefficients in the representation) over the size of the image. For different image and framework sizes, the lower this ratio, the higher the sparseness. As shown in Figure 14.3-B, we indeed see that sparseness increases relative to an increase in image size. This reflects the fact that sparseness is not only local (few edges coexist at one place) but is also spatial (edges are clustered, and most regions are empty). Such a behavior is also observed in V1 of monkeys as the size of the stimulation is increased from a stimulation over only the classical receptive field to 4 times around it (Vinje and Gallant, 2000).

Note that by definition, our representation of edges is invariant to translations, scalings, and rotations in the plane of the image. We also performed the same edge extraction where images from the database were perturbed by adding independent Gaussian noise to each pixel such that signal-to-noise ratio was halved. Qualitative results are degraded but qualitatively similar. In particular, edge extraction in the presence of noise may result in false positives. Quantitatively, one observes that the representation is slightly less sparse. This confirms our intuition that sparseness is causally linked to the efficient extraction of edges in the image.

⁴These python scripts are available at <https://github.com/bicv/SparseEdges> and documented at <https://pythonhosted.org/SparseEdges>.

⁵This database is publicly available at <http://cbcl.mit.edu/software-datasets/serre/SerreOlivaPoggioPNAS07>.

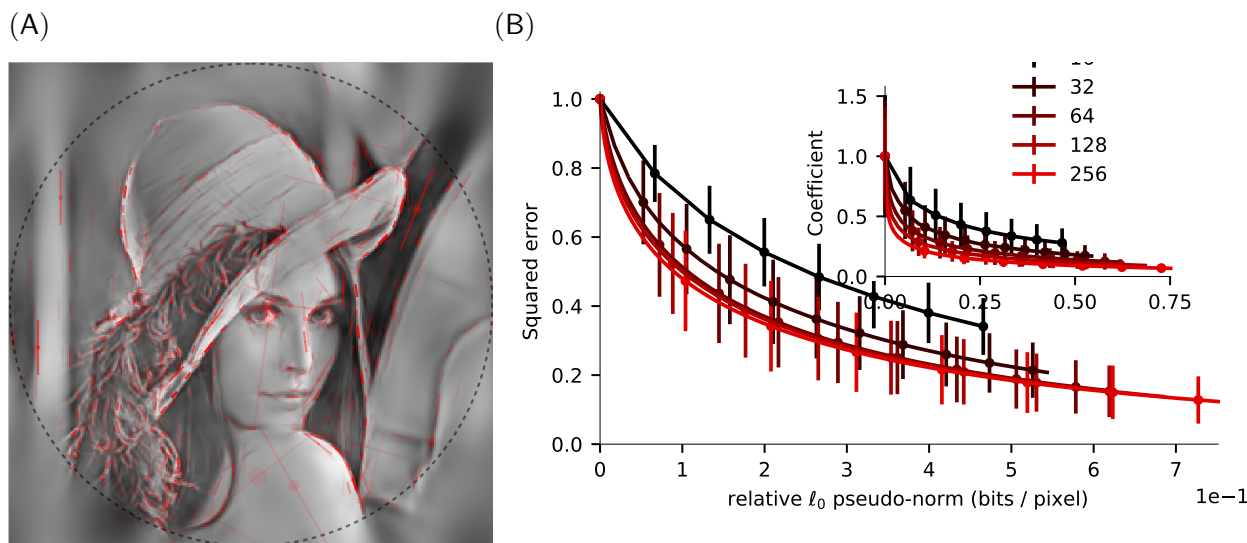


Figure 14.3: **SparseLets framework.** (A) An example reconstructed image with the list of extracted edges overlaid. As in Geisler et al. (2001), edges outside the circle are discarded to avoid artifacts. Parameters for each edge are the position, , orientation, scale (length of bar) and scalar amplitude (transparency) with the phase (hue). We controlled the quality of the reconstruction from the edge information such that the residual energy is less than 3% over the whole set of images, a criterion met on average when identifying 2048 edges per image for images of size 256×256 (that is, a relative sparseness of $\approx 0.01\%$ of activated coefficients). (B) Efficiency for different image sizes as measured by the decrease of the residual’s energy as a function of the coding cost (relative ℓ_0 pseudo-norm). (B, inset) This shows that as the size of images increases, sparseness increases, validating quantitatively our intuition on the sparse positioning of objects in natural images. Note, that the improvement is non significant for a size superior to 128. The SparseLets framework thus shows that sparse models can be extended to full-scale natural images, and that increasing the size improves sparse models by a degree of order (compare a size of 16 with that of 256).

To examine the robustness of the framework and of sparse models in general, we examined how results changed when changing parameters for the algorithm. In particular, we investigated the effect of filter parameters B_f and B_θ . We also investigated how the over-completeness factor could influence the result. We manipulated the number of discretization steps along the spatial frequency axis N_f (that is, the number of layers in the pyramid) and orientation axis N_θ . Results are summarized in Figure 14.4 and show that an optimal efficiency is achieved for certain values of these parameters. These optimal values are in the order of what is found for the range of selectivities observed in V1. Note that these values may change across categories. Further experiments should provide with an adaptation mechanism to allow finding the best parameters in an unsupervised manner.

These particular results illustrate the potential of sparse models in computer vision. Indeed, one main advantage of these methods is to explicitly represent edges. A direct application of sparse models is the ability of the representation to reconstruct these images and therefore to use it for compression (Perrinet et al., 2004). Other possible applications are image filtering or edge manipulation for texture synthesis or denoising (Portilla and Simoncelli, 2000). Recent advances have shown that such representations could be used for the classification of natural images (see chapter 013_theriault or for instance (Perrinet and Bednar, 2015)) or of medical images of emphysema (Nava et al., 2013). Classification was also used in a sparse model for the quantification of artistic style through sparse coding analysis in the drawings of Pieter Bruegel the Elder (Hughes et al., 2010). These examples illustrate the different applications of sparse representations and in the following we will illustrate

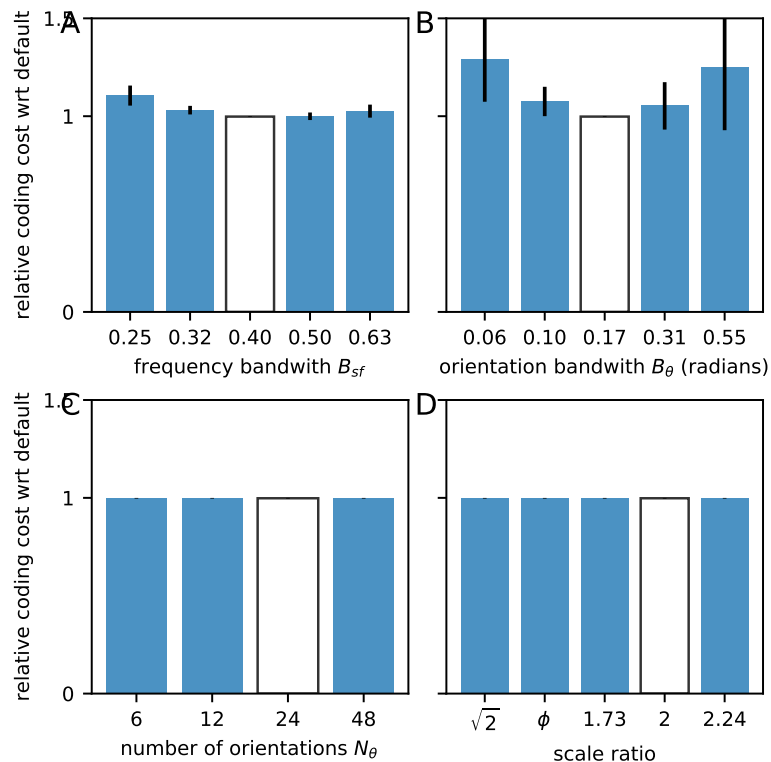


Figure 14.4: **Effect of filters' parameters on the efficiency of the SparseLets framework**

As we tested different parameters for the filters, we measured the gain in efficiency for the algorithm as the ratio of the code length to achieve 85% of energy extraction relative to that for the default parameters (white bar). The average is computed on the same database of natural images and error bars denote the standard deviation of gain over the database. First, we studied the effect of the bandwidth of filters respectively in the (A) spatial frequency and (B) orientation spaces. The minimum is reached for the default parameters: this shows that default parameters provide an optimal compromise between the precision of filters in the frequency and position domains for this database. We may also compare pyramids with different number of filters. Indeed from Equation 14.4, efficiency (in bits) is equal to the number of selected filters times the coding cost for the address of each edge in the pyramid. We plot here the average gain in efficiency which shows an optimal compromise respectively for respectively (C) the number of orientations and (D) the number of spatial frequencies (scales). Note first that with more than 12 directions, the gain remains stable. Note also that a dyadic scale ratio (that is of 2) is efficient but that other solutions —such as using the golden section ϕ — prove to be significantly more efficient, though the average gain is relatively small (inferior to 5%).

some potential perspectives to further improve their representation efficiency.

14.4 SparseEdges: introducing prior information

14.4.1 Using the prior in first-order statistics of edges

In natural images, it has been observed that edges follow some statistical regularities that may be used by the visual system. We will first focus on the most obvious regularity which consists in the anisotropic distribution of orientations in natural images (see chapter 003_holly_gerhard for another qualitative characterization of this anisotropy). Indeed, it has been observed that orientations corresponding to cardinals (that is, to verticals and horizontals) are more likely than other orientations (Ganguli and Simoncelli, 2010; Girshick et al., 2011). This is due to the fact that our point of view is most likely pointing toward the horizon while we stand upright. In addition, gravity shaped our surrounding world around horizontals (mainly the ground) and verticals (such as trees or buildings). Psychophysically, this prior knowledge gives rise to the oblique effect (Keil and Cristóbal, 2000). This is even more striking in images of human scenes (such as a street, or inside a building) as humans mainly build their environment (houses, furnitures) around these cardinal axes. However, we assumed in the cost defined above (see Eq. 14.2) that each coefficient is independently distributed.

It is believed that an homeostasis mechanism allows one to optimize this cost knowing this prior information (Laughlin, 1981; Perrinet, 2010). Basically, the solution is to put more filters where there are more orientations (Ganguli and Simoncelli, 2010) such that coefficients are uniformly distributed. In fact, since neural activity in the assembly actually represents the sparse coefficients, we may understand the role of homeostasis as maximizing the average representation cost $\mathcal{C}(\mathbf{a}|\Phi)$. This is equivalent to saying that homeostasis should act such that at any time, and invariantly to the selectivity of features in the dictionary, the probability of selecting one feature is uniform across the dictionary. This optimal uniformity may be achieved in all generality by using an equalization of the histogram (Atick, 1992). This method may be easily derived if we know the probability distribution function dP_i of variable a_i (see Figure 14.5-A) by choosing a non-linearity as the cumulative distribution function (see Figure 14.5-B) transforming any observed variable \bar{a}_i into:

$$z_i(\bar{a}_i) = P_i(a_i \leq \bar{a}_i) = \int_{-\infty}^{\bar{a}_i} dP_i(a_i) \quad (14.11)$$

This is equivalent to the change of variables which transforms the sparse vector \mathbf{a} to a variable with uniform probability distribution function in $[0, 1]^M$ (see Figure 14.5-C). This equalization process has been observed in the neural activity of a variety of species and is, for instance, perfectly illustrated in the compound eye of the fly's neural response to different levels of contrast (Laughlin, 1981). It may evolve dynamically to slowly adapt to varying changes, for instance to luminance or contrast values, such as when the light diminishes at twilight. Then, we use these point non-linearities z_i to sample orientation space in an optimal fashion (see Figure 14.5-D).

This simple non-parametric homeostatic method is applicable to the SparseLets algorithm by simply using the transformed sampling of the orientation space. It is important to note that the MP algorithm is non-linear and the choice of one element at any step may influence the rest of the choices. In particular, while orientations around cardinals are more prevalent in natural images (see Figure 14.6-A), the output histogram of detected edges is uniform (see Figure 14.6-B). To quantify the gain in efficiency, we measured the residual energy in the SparseLets framework with or without including this prior knowledge. Results show that for a similar number of extracted edges, residual energy is not significantly changed (see Figure 14.6-C). This is again due to the exponential convergence of the squared error (Mallat, 1998, p. 422) on the space spanned by the representation basis. As the tiling of the Fourier space by the set of filters is complete, one is assured of the convergence of the representation in both cases. However thanks to the use of first-order statistics,




Figure 14.5: **Histogram equalization** From the edges extracted in the images from the natural scenes database, we computed sequentially (clockwise, from the bottom left): (A) the histogram and (B) cumulative histogram of edge orientations. This shows that as was reported previously (see for instance (Girshick et al., 2011)), cardinals axis are over-represented. This represents a relative inefficiency as the representation in the SparseLets framework represents *a priori* orientations in an uniform manner. A neuromorphic solution is to use histogram equalization, as was first shown in the fly’s compound eye by (Laughlin, 1981). (C) We draw a uniform set of scores on the y-axis of the cumulative function (black horizontal lines), for which we select the corresponding orientations (red vertical lines). Note that by convention these are wrapped up to fit in the $(-\pi/2, \pi/2]$ range. (D) This new set of orientations is defined such that they are *a priori* selected uniformly. Such transformation was found to well describe a range of psychological observations (Ganguli and Simoncelli, 2010) and we will now apply it to our framework.

the orientation of edges are distributed such as to maximize the entropy, further improving the efficiency of the representation.

This novel improvement to the SparseLets algorithm illustrates the flexibility of the Matching Pursuit framework. This proves that by introducing the prior on first-order statistics, one improves the efficiency of the model for this class of natural images. Of course, this gain is only valid for natural images and would disappear for images where cardinals would not dominate. This is the case for images of close-ups (microscopy) or where gravity is not prevalent such as aerial views. Moreover, this is obviously just a first step as there is more information from natural images that could be taken into account.

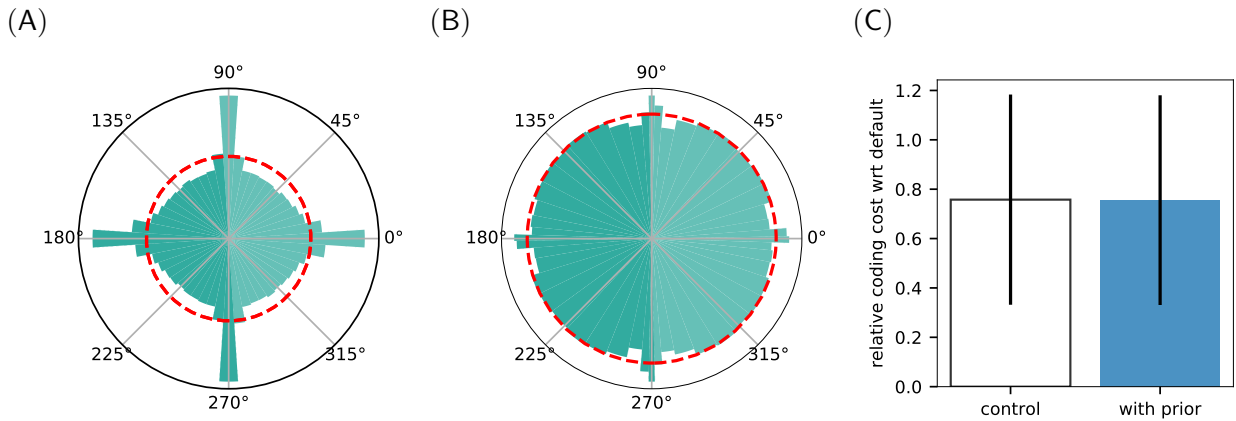


Figure 14.6: **Including prior knowledge on the orientation of edges in natural images.** (A) Statistics of orientations in the database of natural images as shown by a polar bar histogram (the surface of wedges is proportional to occurrences). As for Figure 14.5-A, this shows that orientations around cardinals are more likely than others (the dotted red line shows the average occurrence rate). (B) We show the histogram with the new set of orientations that has been used. Each bin is selected approximately uniformly. Note the variable width of bins. (C) We compare the efficiency of the modified algorithm where the sampling is optimized thanks to histogram equalization described in Figure 14.5 as the average residual energy with respect to the number of edges. This shows that introducing a prior information on the distribution of orientations in the algorithm may also introduce a slight but insignificant improvement in the sparseness.

14.4.2 Using the prior statistics of edge co-occurrences

A natural extension of the previous result is to study the co-occurrences of edges in natural images. Indeed, images are not simply built from independent edges at arbitrary positions and orientations but tend to be organized along smooth contours that follow for instance the shape of objects. In particular, it has been shown that contours are more likely to be organized along co-circular shapes (Sigman et al., 2001). This reflects the fact that in nature, round objects are more likely to appear than random shapes. Such a statistical property of images seems to be used by the visual system as it is observed that edge information is integrated on a local "association field" favoring co-linear or co-circular edges (see chapter 013_theriault section 5 for more details and a mathematical description). In V1 for instance, neurons coding for neighboring positions are organized in a similar fashion. We have previously seen that statistically, neurons coding for collinear edges seem to be anatomically connected (Bosking et al., 1997; Hunt et al., 2011) while rare events (such as perpendicular occurrences) are functionally inhibited (Hunt et al., 2011).

Using the probabilistic formulation of the edge extraction process (see Section 14.2), one can also apply this prior probability to the choice mechanism (Matching) of the Matching Pursuit algorithm. Indeed at any step of the edge extraction process, one can include the knowledge gained by the extraction of previous edges, that is, the set $\mathcal{I} = \{\pi_i\}$ of extracted edges, to refine the log-likelihood of a new possible edge $\pi_* = (*, a_*)$ (where $*$ corresponds to the address of the chosen filter, and therefore to its position, orientation and scale). Knowing the probability of co-occurrences $p(\pi_*|\pi_i)$ from the statistics observed in natural images (see Figure 14.7), we deduce that the cost is now at any coding step (where \mathbf{I} is the residual image —see Equation 14.9):

$$\mathcal{C}(\pi_*|\mathbf{I}, \mathcal{I}) = \frac{1}{2\sigma_n^2} \|\mathbf{I} - a_*\Phi_*\|^2 - \eta \sum_{i \in \mathcal{I}} \|a_i\| \cdot \log p(\pi_*|\pi_i) \quad (14.12)$$

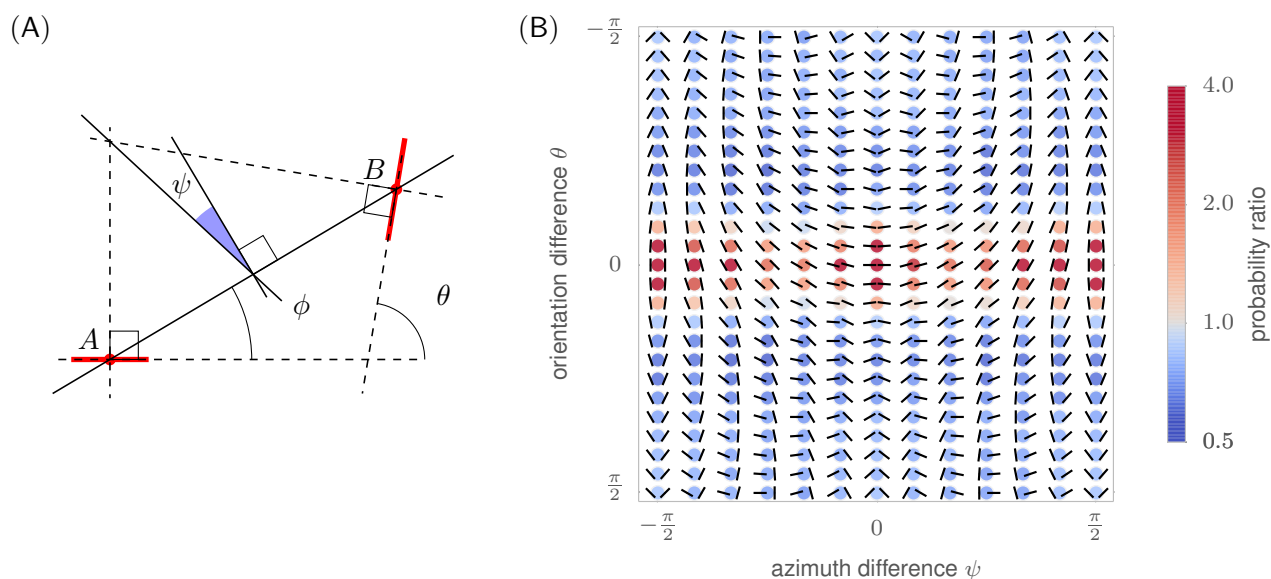


Figure 14.7: **Statistics of edge co-occurrences** (A) The relationship between a pair of edges can be quantified in terms of the difference between their orientations θ , the ratio of scale σ relative to the reference edge, the distance $d = \|\vec{AB}\|$ between their centers, and the difference of azimuth (angular location) ϕ of the second edge relative to the reference edge. Additionally, we define $\psi = \phi - \theta/2$ as it is symmetric with respect to the choice of the reference edge, in particular, $\psi = 0$ for co-circular edges. (B) The probability distribution function $p(\psi, \theta)$ represents the distribution of the different geometrical arrangements of edges' angles, which we call a “chevron map”. We show here the histogram for natural images, illustrating the preference for co-linear edge configurations. For each chevron configuration, deeper and deeper red circles indicate configurations that are more and more likely with respect to a uniform prior, with an average maximum of about 4 times more likely, and deeper and deeper blue circles indicate configurations less likely than a flat prior (with a minimum of about 0.7 times as likely). Conveniently, this “chevron map” shows in one graph that natural images have on average a preference for co-linear and parallel angles (the horizontal middle axis), along with a slight preference for co-circular configurations (middle vertical axis).

where η quantifies the strength of this prediction. Basically, this shows that, similarly to the association field proposed by (Grossberg, 1984) which was subsequently observed in cortical neurons (von der Heydt et al., 1984) and applied by (Field et al., 1993), we facilitate the activity of edges knowing the list of edges that were already extracted. This comes as a complementary local interaction to the inhibitory local interaction implemented in the Pursuit step (see Equation 14.9) and provides a quantitative algorithm to the heuristics proposed in (Fischer et al., 2007a). Note that though this model is purely sequential and feed-forward, this results possibly in a "chain rule" as when edges along a contour are extracted, this activity is facilitated along it as long as the image of this contour exists in the residual image. Such a "chain rule" is similar to what was used to model psychophysical performance (Geisler et al., 2001) or to filter curves in images (August and Zucker, 2001). Our novel implementation provides with a rapid and efficient solution that we illustrate here on a segmentation problem (see Figure 14.8).




Figure 14.8: **Application to rapid contour segmentation** We applied the original sparse edge framework to (A) the synthetic image of a circle embedded in noise. This noise is synthesized by edges at random positions, orientations and scales with a similar first-order statistics as natural scenes. (B) We overlay in red the set of edges which were detected by the SparseLets framework. (C) Then, we have introduced second-order information in the evaluation of the probability in the sparse edges framework (with $\eta = 0.15$). This modified the sequence of extracted edges as shown in blue. There is a clear match of the edge extraction with the circle, as would be predicted by psychophysical results of hand segmentation of contours. This shows that second-order information as introduced in this feed-forward chain may be sufficient to account for contour grouping and may not necessitate a recursive chain rule such as implemented in Geisler et al. (2001).

Indeed, sparse models have been shown to foster numerous applications in computer vision. Among these are algorithms for segmentation in images (Spratling, 2013a) or for classification (Spratling, 2013b; Dumoulin et al., 2014). We may use the previous application of our algorithm to evaluate the probability of edges belonging to the same contour. We show in Figure 14.8 the application of such a formula (in panel C versus classical sparse edge extraction in panel B) on a synthetic image of a circle embedded in noise (panel A). It shows that, while some edges from the background are extracted in the plain SparseLets framework (panel B), edges belonging to the same circular contour pop-out from the computation similarly to a chain rule (panel C). Note that contrary to classical hierarchical models, these results are done with a simple layer of edge detection filters which communicate through local diffusion. An important novelty to note in this extension of the SparseLets framework is that there is no recursive propagation, as the greedy algorithm is applied in a sequential manner.

These types of interaction have been found in area V1. Indeed, the processing may be modulated by simple contextual rules such as favoring co-linear versus co-circular edges (McManus et al., 2011). Such type of modulation opens a wide range of potential applications to computer vision such as robust segmentation and algorithms for the autonomous classification of images (Perrinet and Bednar, 2015). More generally, it shows that simple feed-forward algorithms such as the one we propose may be sufficient to account for the sparse representation of images in lower-level visual areas.

14.5 Conclusion

In this chapter, we have shown sparse models at increasing structural complexities mimicking the key features of the primary visual cortex in primates. By focusing on a concrete example, the SparseLets framework, we have shown that sparse models provide an efficient framework for biologically-inspired computer vision algorithms. In particular, by including contextual information, such as prior statistics on natural images, we could improve the efficiency of the sparseness of the representation.

Such an approach allows to implement a range of practical concepts (such as the good continuity of contours) in a principled way. Indeed, we based our reasoning on inferential processes such as they are reflected in the organization of neural structures. For instance, there is a link between co-circularity and the structure of orientation maps (Hunt et al., 2009). This should be included in further perspectives of these sparse models.

As we saw, the (visual) brain is not a computer. Instead of using a sequential stream of semantic symbols, it uses statistical regularities to derive predictive rules. These computations are not written explicitly, as it suffices that they emerge from the collective behavior in populations of neurons. As such, these rules are massively parallel, asynchronous and error prone. Luckily, such neuromorphic computing architectures begin to emerge — yet, we lack a better understanding of how we may implement computer vision algorithms on such hardware.

As a conclusion, this drives the need for more biologically-driven computer vision algorithm and of a better understanding of V1. However, such knowledge is largely incomplete (Olshausen and Field, 2005) and we need to develop a better understanding of results from electro-physiology. A promising approach in that sense is to include model-driven stimulation of physiological studies (Sanz-Leon et al., 2012; Simoncini et al., 2012) as they systematically test neural computations for a given visual task.

Acknowledgments

The author was supported by EC IP project FP7-269921, “BrainScaleS”. Correspondence and requests for materials should be addressed to the author⁶. Code and supplementary material available at <http://invibe.net/LaurentPerrinet/Publications/Perrinet15bicv>.

⁶Laurent.Perrinet@univ-amu.fr

Bibliography

- Laurent U. Perrinet. *Sparse Models for Computer Vision*, chapter 14. Wiley-VCH Verlag GmbH & Co. KGaA, November 2015. ISBN 9783527680863. doi: 10.1002/9783527680863.ch14. URL <http://onlinelibrary.wiley.com/doi/10.1002/9783527680863.ch14/summary>.
- Karl Friston, Rick A. Adams, Laurent Perrinet, and Michael Breakspear. Perceptions as hypotheses: Saccades as experiments. *Frontiers in psychology*, 3, 2012. ISSN 1664-1078. doi: 10.3389/fpsyg.2012.00151. URL <http://dx.doi.org/10.3389/fpsyg.2012.00151>.
- D. H. Hubel and T. N. Wiesel. Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195(1):215–243, March 1968. ISSN 0022-3751. URL <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1557912/>.
- Alison L. Barth and James F. A. Poulet. Experimental evidence for sparse firing in the neocortex. *Trends in Neurosciences*, 35(6):345–355, June 2012. ISSN 01662236. doi: 10.1016/j.tins.2012.03.008. URL <http://dx.doi.org/10.1016/j.tins.2012.03.008>.
- Ben D. Willmore, James A. Mazer, and Jack L. Gallant. Sparse coding in striate and extrastriate visual cortex. *Journal of neurophysiology*, 105(6):2907–2919, June 2011. ISSN 1522-1598. doi: 10.1152/jn.00594.2010. URL <http://dx.doi.org/10.1152/jn.00594.2010>.
- D. Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Henry Holt & Company, June 1983. URL <http://www.worldcat.org/isbn/0716715678>.
- F. Attneave. Some informational aspects of visual perception. *Psychological Review*, 61(3):183–93, May 1954. ISSN 0033-295X. URL <http://view.ncbi.nlm.nih.gov/pubmed/13167245>.
- Joseph J. Atick. Could information theory provide an ecological theory of sensory processing? *Network: Computation in Neural Systems*, 3(2):213–52, 1992. doi: 10.1088/0954-898X/3/2/009. URL <http://dx.doi.org/10.1088/0954-898X/3/2/009>.
- Jason Wolfe, Arthur R. Houweling, and Michael Brecht. Sparse and powerful cortical spikes. *Current Opinion in Neurobiology*, 20(3):306–312, June 2010. ISSN 09594388. doi: 10.1016/j.conb.2010.03.006. URL <http://dx.doi.org/10.1016/j.conb.2010.03.006>.
- David J. Field. What is the goal of sensory coding? *Neural Computation*, 6(4):559–601, July 1994. ISSN 0899-7667. doi: 10.1162/neco.1994.6.4.559. URL <http://dx.doi.org/10.1162/neco.1994.6.4.559>.
- Emmanouil Froudarakis, Philipp Berens, Alexander S. Ecker, R. James Cotton, Fabian H. Sinz, Dimitri Yatsenko, Peter Saggau, Matthias Bethge, and Andreas S. Tolias. Population code in mouse V1 facilitates readout of natural scenes through increased sparseness. *Nature Neuroscience*, 17(6):851–857, April 2014. ISSN 1546-1726. doi: 10.1038/nn.3707. URL <http://dx.doi.org/10.1038/nn.3707>.

- Günther Palm. Neural associative memories and sparse coding. *Neural Networks*, 37:165–171, January 2013. ISSN 08936080. doi: 10.1016/j.neunet.2012.08.013. URL <http://dx.doi.org/10.1016/j.neunet.2012.08.013>.
- William H. Bosking, Y. Zhang, B. Schofield, and D. Fitzpatrick. Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *Journal of Neuroscience*, 17(6):2112–2127, March 1997. ISSN 1529-2401. URL <http://www.jneurosci.org/cgi/content/abstract/17/6/2112>.
- Jonathan J. Hunt, William H. Bosking, and Geoffrey J. Goodhill. Statistical structure of lateral connections in the primary visual cortex. *Neural Systems & Circuits*, 1(1):3+, 2011. ISSN 2042-1001. doi: 10.1186/2042-1001-1-3. URL <http://dx.doi.org/10.1186/2042-1001-1-3>.
- David J. Field. Relations between the statistics of natural images and the response properties of cortical cells. *Journal of Optical Society of America, A*, 4(12):2379–94, 1987.
- William E. Vinje and Jack L. Gallant. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, 287(5456):1273–1276, February 2000. ISSN 0036-8075. doi: 10.1126/science.287.5456.1273. URL <http://dx.doi.org/10.1126/science.287.5456.1273>.
- Michael R. DeWeese, Michael Wehr, and Anthony M. Zador. Binary spiking in auditory cortex. *Journal of Neuroscience*, 23(21):7940–7949, August 2003. URL <http://www.jneurosci.org/content/23/21/7940.abstract>.
- Pierre Baudot, Manuel Levy, Olivier Marre, Cyril Monier, Marc Pananceau, and Yves Frégnac. Animation of natural scene by virtual eye-movements evokes high precision and low noise in V1 neurons. *Frontiers in Neural Circuits*, 7, 2013. doi: 10.3389/fncir.2013.00206. URL <http://dx.doi.org/10.3389/fncir.2013.00206>.
- Rajesh Rao and Dana Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87, January 1999. ISSN 1097-6256. doi: 10.1038/4580. URL <http://dx.doi.org/10.1038/4580>.
- Michael W. Spratling. A single functional model accounts for the distinct properties of suppression in cortical area V1. *Vision Research*, 51(6):563–576, March 2011. ISSN 1878-5646. doi: 10.1016/j.visres.2011.01.017. URL <http://dx.doi.org/10.1016/j.visres.2011.01.017>.
- Laurent U. Perrinet, Rick A. Adams, and Karl J. Friston. Active inference, eye movements and oculomotor delays. *Biological Cybernetics*, 108(6):777–801, December 2014. ISSN 1432-0770. doi: 10.1007/s00422-014-0620-8. URL <http://dx.doi.org/10.1007/s00422-014-0620-8>.
- Bruno A. Olshausen and David J. Field. Natural image statistics and efficient coding. *Network: Computation in Neural Systems*, 7(6583):333–9, June 1996. ISSN 0028-0836. doi: 10.1038/381607a0. URL <http://dx.doi.org/10.1038/381607a0>.
- Bruno A. Olshausen and David J. Field. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, 37(23):3311–3325, December 1997. ISSN 0042-6989. doi: 10.1016/S0042-6989(97)00169-7. URL [http://dx.doi.org/10.1016/S0042-6989\(97\)00169-7](http://dx.doi.org/10.1016/S0042-6989(97)00169-7).
- Anthony J. Bell and Terrence J. Sejnowski. The ‘independent components’ of natural scenes are edge filters. *Vision Research*, 37(23):3327–38, 1997.
- Colin Fyfe and Roland Baddeley. Finding compact and sparse-distributed representations of visual images. *Network: Computation in Neural Systems*, 6(3):333–344, August 1995. ISSN 1361-6536. doi: 10.1088/0954-898X/6/3/002. URL <http://dx.doi.org/10.1088/0954-898X/6/3/002>.

- Michael Zibulevsky and Barak A. Pearlmutter. Blind Source Separation by sparse decomposition. *Neural Computation*, 13(4):863–82, 2001.
- L. Perrinet, M. Samuelides, and S. Thorpe. Coding static natural images using spiking event times: Do neurons cooperate? *IEEE Transactions on Neural Networks*, 15(5):1164–1175, September 2004. ISSN 1045-9227. doi: 10.1109/TNN.2004.833303. URL <http://dx.doi.org/10.1109/TNN.2004.833303>. Special issue on 'Temporal Coding for Neural Information Processing'.
- Martin Rehn and Friedrich T. Sommer. A model that uses few active neurones to code visual input predicts the diverse shapes of cortical receptive fields. *Journal of Computational Neuroscience*, 22(2):135–46, 2007.
- Eizaburo Doi, Doru C. Balcan, and Michael S. Lewicki. Robust coding over noisy overcomplete channels. *IEEE Transactions in Image Processing*, 16(2):442–52, 2007.
- Laurent U. Perrinet. Role of homeostasis in learning sparse representations. *Neural Computation*, 22(7):1812–1836, July 2010. ISSN 1530-888X. doi: 10.1162/neco.2010.05-08-795. URL <http://invibe.net/LaurentPerrinet/Publications/Perrinet10shl>.
- Michael S. Lewicki and Terrence J. Sejnowski. Learning overcomplete representations. *Neural Computation*, 12(2):337–65, 2000.
- Evan C. Smith and Michael S. Lewicki. Efficient auditory coding. *Nature*, 439(7079):978–982, February 2006. ISSN 1476-4679. doi: 10.1038/nature04485. URL <http://dx.doi.org/10.1038/nature04485>.
- R. Baddeley, L. F. Abbott, M. C. A. Booth, F. Sengpiel, T. Freeman, E. A. Wakeman, and E. T. Rolls. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society B: Biological Sciences*, 264(1389):1775–1783, December 1997. ISSN 1471-2954. doi: 10.1098/rspb.1997.0246. URL <http://citeseer.nj.nec.com/19262.html>.
- Donald O. Hebb. *The organization of behavior: A neuropsychological theory*. Wiley, New York, 1949.
- Joel Zylberberg, Jason T. Murphy, and Michael R. DeWeese. A sparse coding model with synaptically local plasticity and spiking neurons can account for the diverse shapes of V1 simple cell receptive fields. *PLoS Computational Biology*, 7(10):e1002250+, October 2011. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1002250. URL <http://dx.doi.org/10.1371/journal.pcbi.1002250>.
- Jonathan J. Hunt, Peter Dayan, and Geoffrey J. Goodhill. Sparse coding can predict primary visual cortex receptive field changes induced by abnormal visual input. *PLoS Computational Biology*, 9(5):e1003005+, May 2013. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1003005. URL <http://dx.doi.org/10.1371/journal.pcbi.1003005>.
- Florian Blättler and Richard H. R. Hahnloser. An efficient coding hypothesis links sparsity and selectivity of neural responses. *PLoS ONE*, 6(10):e25506+, October 2011. doi: 10.1371/journal.pone.0025506. URL <http://dx.doi.org/10.1371/journal.pone.0025506>.
- Peggy Seriès, Peter E. Latham, and Alexandre Pouget. Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. *Nature Neuroscience*, 7(10):1129–1135, October 2004. ISSN 1097-6256. doi: 10.1038/nn1321. URL <http://dx.doi.org/10.1038/nn1321>.
- Surya Ganguli and Haim Sompolinsky. Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Annual Review of Neuroscience*, 35(1):485–508, 2012. ISSN 1545-4126. doi: 10.1146/annurev-neuro-062111-150410. URL <http://keck.ucsf.edu/~surya/12.CompSense.pdf>.

- Michael Elad. *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer, 1st edition. edition, August 2010. ISBN 144197010X. URL <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/144197010X>.
- A. Benoit, A. Caplier, B. Durette, and J. Hérault. Using human visual system modeling for bio-inspired low level image processing. *Computer Vision and Image Understanding*, 114(7):758–773, July 2010. ISSN 1077-3142. doi: 10.1016/j.cviu.2010.01.011. URL <http://dx.doi.org/10.1016/j.cviu.2010.01.011>.
- Thomas Serre and Tomaso Poggio. A neuromorphic approach to computer vision. *Communications of the ACM*, 53(10):54–61, October 2010. ISSN 0001-0782. doi: 10.1145/1831407.1831425. URL <http://dx.doi.org/10.1145/1831407.1831425>.
- Sylvain Fischer, Rafael Redondo, Laurent U. Perrinet, and Gabriel Cristóbal. Sparse approximation of images inspired from the functional architecture of the primary visual areas. *EURASIP Journal on Advances in Signal Processing*, 2007(1):090727–122, 2007a. ISSN 1687-6180. doi: 10.1155/2007/90727. URL <http://dx.doi.org/10.1155/2007/90727>.
- Laurent U. Perrinet. Adaptive sparse spike coding : applications of neuroscience to the compression of natural images. In Gabriel C. Peter Schelkens, editor, *Optical and Digital Image Processing Conference 7000 - Proceedings of SPIE Volume 7000, 7 - 11 April 2008*, volume 7000. SPIE, 2008.
- Laurent U. Perrinet and Guillaume S. Masson. Modeling spatial integration in the ocular following response using a probabilistic framework. *Journal of Physiology-Paris*, 101(1-3):46–55, January 2007. ISSN 0928-4257. doi: 10.1016/j.jphysparis.2007.10.011. URL <http://dx.doi.org/10.1016/j.jphysparis.2007.10.011>.
- Jorma Rissanen. Modeling by shortest data description. *Automatica*, 14:465–71, 1978.
- Bruno A. Olshausen. Sparse codes and spikes. In Rajesh P. N. Rao, Bruno A. Olshausen, and Michael S. Lewicki, editors, *Probabilistic Models of the Brain: Perception and Neural Function*, chapter Sparse Codes and Spikes, pages 257–72. MIT Press, 2002.
- Andrei N. Tikhonov. *Solutions of Ill-Posed Problems*. Winston & Sons, Washington, 1977. ISBN 0470991240. URL <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0470991240>.
- Matthias Bethge, David Rotermund, and Klaus Pawelzik. Second order phase transition in neural rate coding: Binary encoding is optimal for rapid signal transmission. *Physical Review Letters*, 90(8):088104, 2003.
- Alexander P. Nikitin, Nigel G. Stocks, Robert P. Morse, and Mark D. McDonnell. Neural population coding is optimized by discrete tuning curves. *Physical Review Letters*, 103(13):138101, 2009.
- Hirofugu Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19:716–23, 1974.
- Stéphane Mallat. *A wavelet tour of signal processing*. Academic Press, second edition, 1998.
- Arthur E. C. Pece. The problem of sparse image coding. *Journal of Mathematical Imaging and Vision*, 17:89–108, 2002.
- Daniel D. Lee and Sebastian H. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–91, 1999.

- Marc' A. Ranzato, Christopher S. Poultney, Sumi Chopra, and Yan LeCun. Efficient learning of sparse overcomplete representations with an Energy-Based model. In *Advances in neural information processing systems*, volume 19, pages 1137–44, 2007.
- Jürgen Bolz and Charles D. Gilbert. The role of horizontal connections in generating long receptive fields in the cat visual cortex. *European Journal of Neuroscience*, 1(3):263–8, 1989.
- Erkki Oja. A Simplified Neuron Model as a Principal Component Analyzer. *Journal of Mathematical biology*, 15:267–73, 1982.
- Jiqian Liu and Yunde Jia. Hebbian-based mean shift for learning the diverse shapes of V1 simple cell receptive fields. *Chinese Science Bulletin*, 59(4):452–458, January 2014. ISSN 1861-9541. doi: 10.1007/s11434-013-0041-4. URL <http://dx.doi.org/10.1007/s11434-013-0041-4>.
- Robert Peharz and Franz Pernkopf. Sparse nonnegative matrix factorization with l0-constraints. *Neurocomputing*, 80:38–46, March 2012. ISSN 0925-2312. doi: 10.1016/j.neucom.2011.09.024. URL <http://dx.doi.org/10.1016/j.neucom.2011.09.024>.
- Adam S. Charles, Pierre Garrigues, and Christopher J. Rozell. A common network architecture efficiently implements a variety of Sparsity-Based inference problems. *Neural Computation*, 24(12):3317–3339, December 2012. ISSN 1530-888X. doi: 10.1162/neco_a_-00372. URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.310.2639&rep=rep1&type=pdf>.
- Michal Aharon, Michael Elad, and Alfred M. Bruckstein. On the uniqueness of overcomplete dictionaries, and a practical way to retrieve them. *Linear Algebra and its Applications*, 416(1):48–67, July 2006. ISSN 00243795. doi: 10.1016/j.laa.2005.06.035. URL <http://dx.doi.org/10.1016/j.laa.2005.06.035>.
- John G. Daugman. Two-dimensional spectral analysis of cortical receptive field profiles. *Vision Res.*, 20(10):847–856, 1980. doi: 10.1016/0042-6989(80)90065-6. URL [http://dx.doi.org/10.1016/0042-6989\(80\)90065-6](http://dx.doi.org/10.1016/0042-6989(80)90065-6).
- D. J. Field. Wavelets, vision and the statistics of natural scenes. *Philos T Roy Soc A*, 357(1760): 2527–2542, September 1999. doi: 10.1098/rsta.1999.0446. URL <http://dx.doi.org/10.1098/rsta.1999.0446>.
- Sylvain Fischer, Filip Sroubek, Laurent U. Perrinet, Rafael Redondo, and Gabriel Cristóbal. Self-invertible 2D log-Gabor wavelets. *International Journal of Computer Vision*, 75(2):231–246, January 2007b. ISSN 1573-1405. doi: 10.1007/s11263-006-0026-8. URL <http://dx.doi.org/10.1007/s11263-006-0026-8>.
- Paul D. King, Joel Zylberberg, and Michael R. DeWeese. Inhibitory interneurons decorrelate excitatory cells to drive sparse code formation in a spiking model of V1. *Journal of Neuroscience*, 33(13):5475–5485, March 2013. ISSN 1529-2401. doi: 10.1523/jneurosci.4188-12.2013. URL <http://dx.doi.org/10.1523/jneurosci.4188-12.2013>.
- Stéphane Mallat and Zhifeng Zhang. Matching Pursuit with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3414, 1993.
- T. E. Oliphant. Python for scientific computing. *Computing in Science and Engineering*, 9(3):10–20, May 2007. ISSN 1521-9615. doi: 10.1109/MCSE.2007.58. URL <http://dx.doi.org/10.1109/MCSE.2007.58>.

- John D. Hunter. Matplotlib: A 2D graphics environment. *Computing in Science and Engineering*, 9(3):90–95, May 2007. ISSN 1521-9615. doi: 10.1109/MCSE.2007.55. URL <http://dx.doi.org/10.1109/MCSE.2007.55>.
- Wilson S. Geisler, J. S. Perry, B. J. Super, and D. P. Gallogly. Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41(6):711–24, March 2001. ISSN 0042-6989. doi: 10.1016/s0042-6989(00)00277-7. URL [http://dx.doi.org/10.1016/s0042-6989\(00\)00277-7](http://dx.doi.org/10.1016/s0042-6989(00)00277-7).
- Javier Portilla and Eero P. Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *Int. J. Comput. Vision*, 40(1):49–70, October 2000. ISSN 0920-5691. doi: 10.1023/a:1026553619983. URL <http://dx.doi.org/10.1023/a:1026553619983>.
- Laurent U. Perrinet and James A. Bednar. Edge co-occurrences can account for rapid categorization of natural versus animal images. *Scientific Reports*, 5:11400, 2015. doi: 10.1038/srep11400. URL <http://www.nature.com/articles/srep11400>. In print.
- Rodrigo Nava, J. Victor Marcos, Boris Escalante-Ramírez, Gabriel Cristóbal, Laurent U. Perrinet, and Raúl S. J. Estépar. *Advances in Texture Analysis for Emphysema Classification*, volume 8259 of *Lecture Notes in Computer Science*, chapter 27, pages 214–221. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013. ISBN 978-3-642-41826-6. doi: 10.1007/978-3-642-41827-3_27. URL http://dx.doi.org/10.1007/978-3-642-41827-3_27.
- J. M. Hughes, D. J. Graham, and D. N. Rockmore. Quantification of artistic style through sparse coding analysis in the drawings of Pieter Bruegel the Elder. *Proceedings of the National Academy of Sciences*, 107(4):1279–1283, January 2010. ISSN 1091-6490. doi: 10.1073/pnas.0910530107. URL <http://dx.doi.org/10.1073/pnas.0910530107>.
- Deep Ganguli and Eero Simoncelli. Implicit encoding of prior probabilities in optimal neural populations. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 658–666. 2010. URL <http://www.cns.nyu.edu/pub/lcv/ganguli10c-preprint.pdf>.
- Ahna R. Girshick, Michael S. Landy, and Eero P. Simoncelli. Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, 14(7):926–932, July 2011. ISSN 1546-1726. doi: 10.1038/nn.2831. URL <http://dx.doi.org/10.1038/nn.2831>.
- Matthias S. Keil and Gabriel Cristóbal. Separating the chaff from the wheat: possible origins of the oblique effect. *Journal of the Optical Society of America A*, 17(4):697+, 2000. ISSN 1084-7529. doi: 10.1364/josaa.17.000697. URL <http://dx.doi.org/10.1364/josaa.17.000697>.
- S. Laughlin. A simple coding procedure enhances a neuron’s information capacity. *Zeitschrift für Naturforschung. Section C: Biosciences*, 36(9-10):910–912, 1981. ISSN 0341-0382. URL <http://view.ncbi.nlm.nih.gov/pubmed/7303823>.
- Mariano Sigman, Guillermo A. Cecchi, Charles D. Gilbert, and Marcelo O. Magnasco. On a common circle: Natural scenes and Gestalt rules. *Proceedings of the National Academy of Sciences*, 98(4):1935–1940, February 2001. ISSN 1091-6490. doi: 10.1073/pnas.031571498. URL <http://dx.doi.org/10.1073/pnas.031571498>.
- Stephen Grossberg. *Outline of A Theory of Brightness, Color, and form Perception*, volume 20, pages 59–86. Elsevier, 1984. ISBN 9780444875129. doi: 10.1016/s0166-4115(08)62080-4. URL [http://dx.doi.org/10.1016/s0166-4115\(08\)62080-4](http://dx.doi.org/10.1016/s0166-4115(08)62080-4).

- R. von der Heydt, E. Peterhans, and G. Baumgartner. Illusory contours and cortical neuron responses. *Science (New York, N. Y.)*, 224(4654):1260–1262, June 1984. ISSN 0036-8075. doi: 10.1126/science.6539501. URL <http://dx.doi.org/10.1126/science.6539501>.
- David J. Field, A. Hayes, and R. F. Hess. Contour integration by the human visual system: evidence for a local "association field". *Vision Research*, 33(2):173–193, January 1993. ISSN 0042-6989. doi: 10.1016/0042-6989(93)90156-Q. URL <http://view.ncbi.nlm.nih.gov/pubmed/8447091>.
- Jonas August and StevenW Zucker. A markov process using curvature for filtering curve images. In Mário Figueiredo, Josiane Zerubia, and AnilK Jain, editors, *Energy Minimization Methods in Computer Vision and Pattern Recognition*, volume 2134 of *Lecture Notes in Computer Science*, pages 497–512. Springer Berlin Heidelberg, 2001. doi: 10.1007/3-540-44745-8_33. URL http://dx.doi.org/10.1007/3-540-44745-8_33.
- Michael W. Spratling. Image segmentation using a sparse coding model of cortical area V1. *IEEE Transactions on Image Processing*, 22(4):1631–1643, April 2013a. ISSN 1941-0042. doi: 10.1109/tip.2012.2235850. URL <http://dx.doi.org/10.1109/tip.2012.2235850>.
- M. W. Spratling. Classification using sparse representations: a biologically plausible approach. *Biological Cybernetics*, 108(1):61–73, December 2013b. ISSN 1432-0770. doi: 10.1007/s00422-013-0579-x. URL <http://dx.doi.org/10.1007/s00422-013-0579-x>.
- Serge O. Dumoulin, Robert F. Hess, Keith A. May, Ben M. Harvey, Bas Rokers, and Martijn Barendregt. Contour extracting networks in early extrastriate cortex. *Journal of Vision*, 14(5):18+, May 2014. ISSN 1534-7362. doi: 10.1167/14.5.18. URL <http://dx.doi.org/10.1167/14.5.18>.
- J. N. J. McManus, Wu Li, and Charles D. Gilbert. Adaptive shape processing in primary visual cortex. *Proceedings of the National Academy of Sciences*, 108(24):9739–9746, May 2011. ISSN 1091-6490. doi: 10.1073/pnas.1105855108. URL <http://dx.doi.org/10.1073/pnas.1105855108>.
- Jonathan J. Hunt, Clare E. Giacomantonio, Hua jin Tang, Duncan Mortimer, Sajjida Jaffer, Vasily Vorobyov, Geoffery Ericksson, Frank Sengpiel, and Geoffrey J. Goodhill. Natural scene statistics and the structure of orientation maps in the visual cortex. *NeuroImage*, 47(1):157–172, August 2009. ISSN 1095-9572. doi: 10.1016/j.neuroimage.2009.03.052. URL <http://dx.doi.org/10.1016/j.neuroimage.2009.03.052>.
- Bruno A. Olshausen and David J. Field. How close are we to understanding V1? *Neural Computation*, 17(8):1665–1699, August 2005. ISSN 1530-888X. doi: 10.1162/0899766054026639. URL <http://portal.acm.org/citation.cfm?id=1118017>.
- Paula Sanz-Leon, I. Vanzetta, G. S. Masson, and L. U. Perrinet. Motion Clouds: model-based stimulus synthesis of natural-like random textures for the study of motion perception. *Journal of Neurophysiology*, 107(11):3217–3226, March 2012. ISSN 1522-1598. doi: 10.1152/jn.00737.2011. URL <http://dx.doi.org/10.1152/jn.00737.2011>.
- Claudio Simoncini, Laurent U. Perrinet, Anna Montagnini, Pascal Mamassian, and Guillaume S. Masson. More is not always better: adaptive gain control explains dissociation between perception and action. *Nat Neurosci*, 15(11):1596–1603, November 2012. ISSN 1097-6256. doi: 10.1038/nn.3229. URL <http://dx.doi.org/10.1038/nn.3229>.

