Computational Explorations in Perceptual Learning

Guillaume Hennequin

Master of Science

Cognitive Science and Natural Language Processing

School of Informatics

University of Edinburgh

2007

Abstract

Performance in a perceptual task improves with practice, a phenomenon known as perceptual learning. Its behavioral aspect has been receiving massive interest for over a century, while what happens in the brain that underlies the usually dramatic perceptual improvement remains largely unknown. Over the last two decades, the psychophysics literature has been focusing on the properties of this learning that can help pinning down a locus of learning. Early areas in the visual pathway seem to be excellent candidates. Only recently, physiologists have been applying recording techniques to look into the neural correlates of perceptual learning in monkey, showing that fine modifications of the stimulus representation occur in those areas.

First, we show that these patterns of neural changes are very well accounted for in a model of unsupervised Hebbian learning relying on the prior distribution of input stimulus.

Second, we use a population coding/decoding approach to relate physiology to perception, and show that not much improvement can actually be achieved on the basis of such representational changes. We discuss alternative options.

Acknowledgements

I am very grateful to Dr. Peggy Seriès, lecturer at the Institute of Adaptative and Neural Computation, for her careful supervision in terms of both scientific content and research skills. I would like to thank Dr. James A. Bednar for providing good advice related to the LISSOM model, as well as Drs. Andrew Teich and Ning Qian for making the matlab code of their 2003 paper [65] available to us.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Guillaume Hennequin)

Table of Contents

Chapter 1

Introduction

Animals and human beings are constantly receiving sensory information from their environment. They often have to make decisions based on what they have just smelt, felt, seen or heard. Because our brain has been evolving over tens of thousands of years, it may be reasonable to think that it has reached some optimality with respect to the tasks we have to perform very often. However, intriguing experiments show that when we train intensively in a perceptual task, we can get better and considerably improve our decisions. This is what is called "perceptual learning" in its broad sense. It shows that what our brain does – from perceiving to deciding – is not naturally optimal, and that "practice makes [it] perfect" [33].

This would not be surprising if it happened only for high level perceptual skills, such as learning to speak, to read, to walk, for which everyone knows that newborn babies are, on the contrary, far from being optimal. Still, perceptual learning has been observed in adults, and for tasks involving low-level stimulus attributes, such as position, orientation, or frequency. In this dissertation, we are mainly interested in learning to discriminate between two differently oriented visual stimuli. Human beings can improve their discrimination threshold by up to 70%. The question of what changes in the brain that yields such a behavioral improvement is still a hot and much debated question. And it is a very challenging one. Many researchers have studied either the psychophysical or the neurophysiological aspect of perceptual learning, but few people have tried to reconcile both sides: can the observed neural changes account for the behavioral improvement?

We start this dissertation by reviewing what is known about the neurophysiology and the psychophysics of perceptual learning (chapter 2).

We then present the framework of population coding (chapter 3), a starting point to all the models we use in this dissertation. It describes how the information about a sensory stimulus is encoded in large pools of cells.

In chapter 4, we first address the question of *how the observed neural changes can be accounted for*, by suggesting and discussing models of learning "without teacher". These models are mainly based on the idea that during the course of learning a task, one is presented the same stimulus repeatedly during hours. The learning algorithm is therefore likely to be very sensitive to the probability distribution of the stimulus. Furthermore, learning has been shown to occur even when the subject is given no feedback [15, 12, 33]. We find that our models can account for the neural changes observed in physiology, and can also interestingly replicate the changes obtained in another model by Teich & Qian, 2003 using completely different methods [65].

As we will see in chapter 2, most psychophysical studies lead to the conclusion of an early locus for the underlying neural correlates: V1, V2 ? Does physiology really agree with that ? In chapter 6, we relate the modifications of neuronal properties to the observed psychophysical performance in the task. More specifically, we address the question of *how much* behavioral improvement is predicted by the neural changes. To this end, we add to our encoding model (chapter 3) a model of stimulus reconstruction (decoding), on the top of which we finally have a model of perceptual decision (chapter 5). This general architecture can be found in figure 5.2, p. 56. By considering only changes of the encoder, we show that physiology can predict *less than half* of the perceptual improvement observed in psychophysics (chapter 6). We conclude that learning the encoder is very likely not to be the full story. We finally discuss the extent to which learning the decoder could account for the rest of the learning.

Chapter 2

Background

2.1 Introduction

As stated in the introduction, "perceptual learning" designates a modification of decision skills related to perception of sensory signals. Learning is in fact a major property of our neural circuitry. Without plasticity, our brain would not be able to evolve from birth, and the reader would certainly not be so well named. Learning has been shaping us for long, and of course this process also crucially depends on our environment, on what we perceive. Therefore, it may seem difficult to extract a clear definition of perceptual learning. Is learning the understanding of a spoken foreign language part of the perceptual learning phenomena? In fact, in the literature, perceptual learning is restricted to very low-level tasks: learning the discrimination of visual contrasts, learning the estimation of the direction of an object's motion, learning to recognize textures, faces, ... To provide a more formal definition of perceptual learning, let us quote [22]:

> *Any relatively permanent and consistent change in the perception of a stimulus array following practice or experience with this array will be considered perceptual learning.*

Another criterion that goes the same way is that of the *non-consciousness* of the learning process. Perceptual learning refers to a form of learning where the subject eventually does not "know" anything new, explicitely. This criterion naturally restricts the definition of perceptual learning to a narrower range of tasks, often involving low-level characteristics of the stimulus in play.

Finally, perceptual learning is restricted to *long-lasting* modifications of the behavior, contrary to the "adaptation" phenomenon which stands for perceptual changes at a smaller time scale and sets the stage for famous visual illusions (see [7] for a review).

Perceptual learning has been observed in all modalities: vision, olfaction, somatosensation, audition. As we shall see later (section 2.3, page 10), the neural correlates of the behavioral improvement are better revealed in the last three modalities than they are for vision. Changes in the brain are dramatic in the auditory, olfactory and somatosensory cortices, whereas they seem to be more subtle in the primary visual cortex.

Physiological studies that we sum up in section 2.3 show that the neural representations of basic features of the visual stimulus change during the course of learning. At the same time, resulting behavioral changes are easily assessed. This makes of perceptual learning a very interesting paradigm to understand the neural code.

In this chapter, we review the psychophysical aspects of perceptual learning – with special emphasis on orientation discrimination – before summing up the recent discoveries related to its neural correlates. The purpose of this introductory chapter is not to provide an exhaustive review of perceptual learning (which is to be found, up to 2002, in [16], or in more recent reviews like [13, 14, 19, 31]) but rather to give a general feeling of the complexity of the phenomenon, as well as a sense of the crucial questions in play, which will prove useful to the understanding of further discussions in this dissertation. Some details about perceptual learning in orientation discrimination are also needed, because a whole chapter herein (chapter 6) is aimed at bridging psychophysics and physiology within this paradigm.

2.2 Psychophysics

Perceptual learning has been studied psychophysically for decades [22]. In such studies, a subject is trained to perform a task under a certain set of conditions, until it reaches an asymptotically stable level of performance. Performance usually increases under this set of conditions as the subjects trains. The transfer of

the learning to other untrained sets of conditions is then tested. Alternatively, the same conditions are used, and the transfer to a new task is evaluated.

2.2.1 A wide range of tasks

Perceptual learning occurs in many submodalities in each sense. It has been probed in a great variety of tasks. Most tasks involve low-level characteristics of the stimulus. In audition, discrimination of sound frequency improves with training [51]. In somatosensation, improvements in tactile discrimination of position [53] and frequency [50] have been reported. Vision offers a much better control of the stimulus compared to other modalities, and has therefore received even greater interest: visual discrimination of position [9], hyperacuity [15, 27], orientation discrimination [54], direction of motion [3, 4] and many more.

Perceptual learning has also been studied in more complex tasks: discrimination between novel faces in noisy contexts [25], research of a particular shape in an array of distractors [61], or identification of its "odd element" [62]. A recent study showed that even training in action video games is capable of altering a range of visual skills [59].

2.2.2 Two different time scales

Learning has been observed to occur at two different time scales (see [34] for a review). A fast learning is often reported within training sessions, and is a matter of minutes or hours. Usually, the resulting increase of performance does not last for long, and the subject has to start learning almost from scratch at the beginning of the next session. "Almost" means that performance does actually increase from session to session, following a lower time constant. This slower form of learning seems to require "consolidation phases", e.g. the night between two training days [54]. In many studies, both time constants have been observed within the same subjects.

A. Karni and G. Bertini [32] speculate that fast learning would reflect "the setting up of a task-specific processing routine for solving the perceptual problem". On the contrary, they think of slow learning as revealing an "ongoing, perhaps structural, modification of basic representations within the processing system".

2.2.3 Specificity of the learning

Learning has proved to be often specific to i) the properties of the stimulus, and ii) the task itself. As discussed below, these two properties of perceptual learning have major implications for our understanding of i) where learning occurs and ii) what is actually being learnt.

Stimulus specificity

Most psychophysical studies of perceptual learning over the last decade have been reporting that the improvement obtained after training is often restricted to stimuli similar to the trained stimulus. For example, in visual tasks where the stimulus involves a given orientation (grating waveform discrimination, texture discrimination, vernier, ...), learning does not transfer to the orthogonal orientation. Even more surprising is the finding in [54, 60] that performance is significantly worse after than before, at the orthogonal orientation (in contradiction with [67] in which an improvement is reported for the same task).

This specificity has been invoked to infer the locus of learning. If learning is specific to a given property of the stimulus, areas where neurons are selective to this property, and better, areas where this property is represented with the finest resolution, are most likely to host the underlying neural modifications. For example, the fact that orientation discrimination does not improve in untrained locations in the visual field gives priority to areas where the retinotopy is preserved, and preferentially where neurons have small receptive fields. Similarly, the specificity to the orientation of the stimulus favors areas where neurons are orientation selective (see 2.3.2 for more details about orientation tuning).

Recently, however, [36] found that learning the discrimination of motion direction can transfer to an untrained direction, provided the difficulty of the task is moderately reduced. They hypothesize that the restriction of the learning may stem from the extreme difficulty of the tasks usually reported in the literature. They conclude from the observed generalization that neural correlates should also be found in higher-level areas even in “low-level” tasks.

Task specificity

In most cases, training in one particular task does not yield any improvement in other tasks involving the same stimulus characteristics. For example, [9] trained observers in a bisection task, where subjects were asked to report whether, in a pattern of three parallel segments |||, the middle segment is shifted toward the left one or toward the right one. They showed that the subjects dramatically improved their levels of performance in this bisection task, whereas they did not improve in an orientation discrimination task involving the same stimulus orientation and position as in the trained bisection task. Similarly, no improvement was found in a vernier discrimination with similar attributes. On the contrary, testing the same observers in the same bisection task, in which the left and right segments were more spaced out, revealed a significant amount of learning.

If the stimulus specificity of perceptual learning has been used to advocate in favor of early neuronal modifications, on the contrary the task specificity suggests that learning may in fact occur at multiple levels of visual processing [12, 9]. Indeed, if the neural basis was to be found in early areas only, tasks involving the same stimulus attributes (finely represented in those areas) would show at least partial transfer of learning from one another.

In some cases, however, learning can transfer to another task. A study by N. Matthews showed that subjects trained in a motion direction discrimination task show better performance in orientation discrimination, (the inverse transfer is not true, though) [39]. The same researcher demonstrated that the improvement in orientation discrimination through training also yields better contrast sensitivity [40].

Studying the task transferability also helps determining what the subjects are actually learning. For example, the above mentioned subjects in [9] certainly develop a strategy specific to a bisection task, that is not helpful in vernier or orientation discrimination. Furthermore, the transfer to other tasks is subject to variability across observers. This also suggests that there may not be one single strategy to solve a perceptual task, but different possibilities, which makes the understanding of the neural basis of perceptual learning all the more challenging.

2.2.4 Subsequent theories

In an attempt to explain all these psychophysical findings (timing, specificity, variability) within clean and unified frameworks, a few theories of the possible learning mechanisms have emerged recently. They are reviewed in [14]. The *reverse hierarchy theory* [1] hypothesizes that high cortical areas are the first ones to undertake modifications, and that lower areas carry on if necessary. In [13], M. Fahle contrasts two theories of signal detection improvement: an *early selection theory* stipulating that learning a perceptual skill is about getting rid of irrelevant signals and raising the signal-to-noise ratio as soon as possible in the visual pathway, and a *late selection theory*, in which irrelevant signals are removed in later stages of cortical information processing, relevant signals being "selected" at those stages. Analogies with attentional processes have been also inspiring these different alternatives.

2.2.5 Orientation discrimination

In this dissertation, orientation discrimination is a case of special interest. We here review a couple of studies, summing up what is currently known of the behavioral aspects of learning this type of discriminative skill.

Learning the discrimination between two slightly different orientations can be achieved by different task paradigms. The one-interval setting consists in showing only one oriented stimuli and asking the observer to tell whether it was rotated clockwise or anticlockwise from a reference which is not shown [54, 49]. In the two-intervals setting, two oriented stimuli with slightly different angles are shown successively. The subject is to tell whether the second stimulus was rotated clockwise or anticlockwise from the first one [56]. Alternatively, one can ask the subject to tell whether the stimuli were tilted in the same direction from an unshown reference [68]. Usually, auditory feedback (correct, incorrect) is provided after each trial.

Different sorts of visual stimuli have been used. Some of them are reported in figure 2.1. The first pattern is a simple oriented line. The second pattern has been designed to avoid giving the subject contextual cues other than orientation difference (luminance change at a certain screen area, rotation illusion, ...).

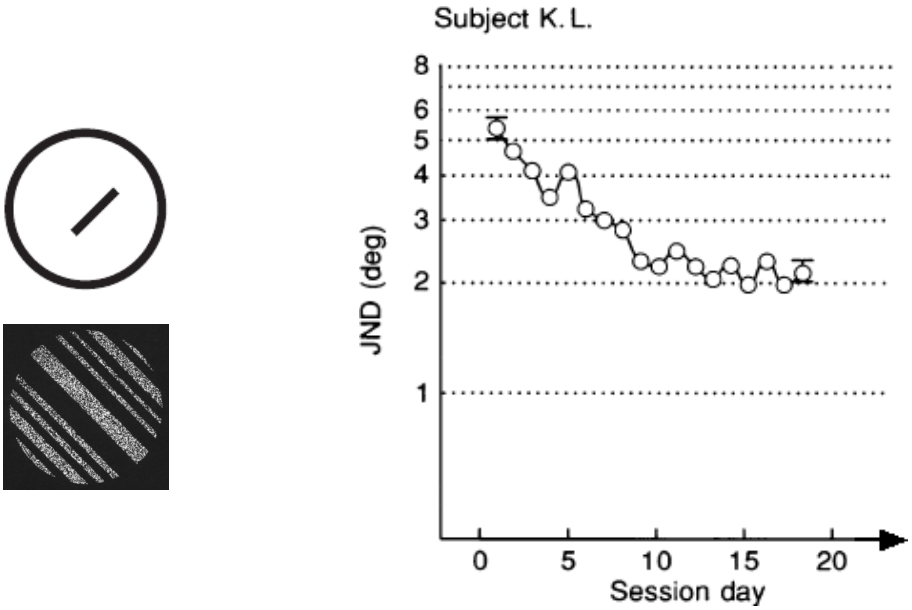

Figure 2.1: **Left**: *different stimuli used in orientation discrimination tasks: a simple oriented line [67, 60, 41, 39] and a noisy grating with random spatial phase [56].* **Right**: *the discrimination performance increases (the Just Noticeable Difference (JND) – see page 59 – decreases) as the subject trains. Reprinted from [54].*

In the domain of orientation discrimination, three influential studies are always quoted: Vogels and Orban 1985 [67], Shiu and Pashler 1992 [60] and Schoups *et al.* 1995 [54]. All studies showed a dramatic improvement after weeks of training (figure 2.1, right). [60] showed that the improvement is restricted to the trained location: when the stimulus is spatially shifted away from the trained location in the visual field, subjects exhibit the same discrimination threshold as before learning, and have to learn from scratch if they want to improve at that new location. The transfer was actually tested for extreme positions in the visual field: improvements for a stimulus positioned in one corner did not transfer to other corners. As mentioned above (section 2.2.3, page 6), this position specificity suggests that the underlying neural changes are more likely to be located in areas where the retinotopy is conserved. However, according to Schoups *et al.*, "the precision of the retinotopy, rather than the retinotopy itself, will be relevant in the localization of the learning effect". Indeed, the retina is topographically mapped into many cortical areas in the visual pathway: from V1 to up to area TEO [6]. Furthermore, the farther the area is from the retina, the worse the precision of the mapping. Thus, in order to pin down one of those areas, one

would need to compare i) the minimal distance from the trained position at which performance drops and ii) the radius of the receptive fields in the area. In light of these considerations, [54] investigated the transfer to new positions using a finer resolution. They showed that the interstimuli distance can be as small as the diameter of the stimuli – stimuli abutting but not overlapping – for the learning not to transfer. They also reported that "the trained stimulus could be located as close as 1.6 deg from the vertical meridian without affecting the other hemifield". They concluded that early areas such as V1 or V2 are most likely to host the neural changes, while V3 and V4 are certainly excluded.

The same study [54] also looked into the interocularity of the learning. They found that training with one eye equally (or almost equally) improves performance of both eyes. They argue that it is not in contradiction with an early locus of learning. The first reason is that binocular cells are activated even when the subject closes one eye, which means that potential learning effects may not necessarily be restricted to monocular cells. The second reason it that although some neurons exhibit strong occular dominance in V1 and V2, this feature seems to be negatively correlated with orientation selectivity [5].

Finally, all 3 mentioned studies showed that the improvement is restricted to the trained orientation (see above, page 6). It should be mentioned, however, that they do not assess the transfer to the full range of orientations. Rather, they only look at what happens at the orientation orthogonal to the trained orientation, assuming that since it is the farthest orientation, intermediate angles would not show less transfer. When looking at the relationships between neural changes and behavioral performance (see chapter 6, page 61), we found that the psychophysics literature is currently lacking a test of transfer across the full range of orientations.

2.3 Neurophysiology

As reviewed in [24], the neural substrate of perceptual learning includes changes in cortical maps, in the temporal characteristics of neuronal responses, and in modulation of contextual influences. The extent of these neural changes differ between sensory modalities. First, we briefly review what happens in the auditory and somatosensory cortices following practice. Second, we sum up a couple

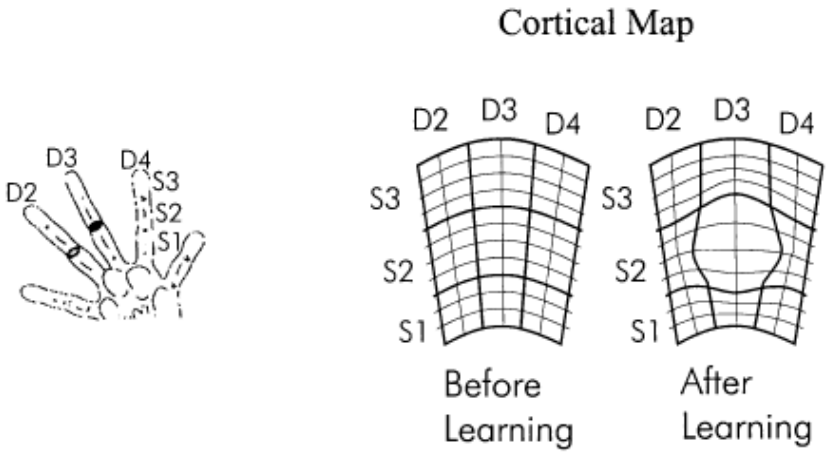

Figure 2.2: *The reorganization of the spatial mapping of the skin surface onto the somatosensory cortex, following intensive practice in a tactile discrimination task. The dark patch on the schematic hand (left) represents the trained area. The cortical map is topographic, hence the direct correpondance between the different parts of the hand and the cortical grids. Reprinted from [50].*

of physiological studies aimed at understanding the neural basis of orientation discrimination learning.

2.3.1 Dramatic reorganization of the cortex

Several studies by G. H. Recanzone have focused on the neural changes in the auditory and somatosensory cortices following intensive practice in sensory tasks [50, 51]. Their results show that these cortices are capable of large scale reorganization in a way that tends to over-represent some attributes of the trained stimulus. For example, [51] shows that, when a monkey is trained in a tactile discrimination task, the area of cortex onto which the trained area of skin on the index finger is mapped extends significantly. Thus, this skin area is over-represented as compared to before learning (see figure 2.2). In the auditory cortex, the representation of the trained stimulus is also made better by a densification of the tonotopic map around the trained frequency.

This kind of topological reorganization of cortical maps has been observed even without the need of running the subject through a supervised perceptual task.

For example, [47] applied a coactivation protocol in which an small area on the index finger of the subject was stimulated during 3 hours – the subject was not attending the stimulation, and behaved as in a normal day of work. The stimulation spanned all the (partially overlapping) receptive fields in the area, such that all neurons were coactivated in a Hebbian manner, which strenghtened their mutual interconnectedness. Before and after this coactivation process, the subject was tested in a spatial tactile discrimination task: his finger was stimulated at two sligthly different positions, and he was to tell whether he felt one or two stimulations. The performance in the task increased after the coactivation period, that is, the distance threshold decreased significantly.

2.3.2 Fine-retuning in the domain of vision

Until recent years, the primary visual cortex was considered to be a fixed "hard-wired" module of visual processing. This was largely supported by Hubel and Wiesel's experiments on the development of cat and monkey visual cortices, in which they showed that if the visual cortex in young animal is still very plastic, no dramatic reorganisation of it can occur after a certain age [29]. Recent studies have changed this view, providing evidence for adult plasticity in these areas (see [23] for a review).

Especially, four recent neurophysiological studies have revealed that the dramatic changes that occur in the somatosensory or auditory cortices are not reproduced in the visual cortex, and that neural correlates of perceptual learning in this locus, if any, are made of more subtle changes of the neuronal tuning properties. We here review what these tuning properties are, before summing up the physiological reports of the four studies.

Tuning properties of visual cortical cells

In V1 (the first cortical area after the retina and the lateral geniculate nucleus (LGN) in the early visual pathway), neurons have receptive fields covering small continuous areas on the retina. Moreover, they are spatially organised such that neighboring neurons in V1 have neighboring receptive fields on the retina ("retinotopy"). As a result, the retina is topographically "mapped" onto the cortex. This

property propagates to further areas in the visual pathway: V2, V4, MT, ... Receptive fields of V1 neurons are smaller than in V2 and V4, where the spatial overlap is therefore greater.

The visual cortex also shows a functional organization in "columns". In one column, hundreds of nearby neurons have similar "preferences" with respect to certain properties of the visual stimulus spanning their receptive field. Thus, in addition of being selective to spatial location due to their receptive fields, neurons in V1, V2 and V4 are selective to contrast, color, orientation, spatial frequency, or exhibit a preference for one of the two eyes (occular dominance). In MT, one find neurons selective to the direction of motion, and even to faces. Preference to those properties is changing smoothly from one column to its neighbors, resulting in a spatial organization of selectivity called a *cortical feature map*. Optical imaging techniques have allowed the measurement of such maps, one of which is given as an example in figure 2.3.

In this dissertation, we are principally concerned by orientation selectivity. Most neurons in the visual cortex respond maximally when the stimulus covering their receptive field has a particular orientation. The response gradually decreases as the stimulus angle steps away from the preferred orientation. Furthermore, physiological recordings show that this response is inherently noisy. The average activity in response to the full range of orientations is called an *orientation tuning curve*. An example of such curve is reported in figure 2.3.

Note that this representation of the neuronal response properties assumes a "rate code", in which the information is encoded in firing rates (measured with a certain time window) while the precise spike timing is not taken into account.

One-neurons studies

Several studies have provided evidence that under certain conditions, the behavioral performance as well as its improvement with practice can be accounted for by the activity of *single neurons*. For example, E. Zohary and coworkers recorded the activity of single cells before and after training a monkey in a motion direction discrimination task. They also combined the recordings with psychophysical experiments in order to get the perceptual performance and improvement. From the recordings, they computed the "neuronal sensitivity", that is, the hypothe-

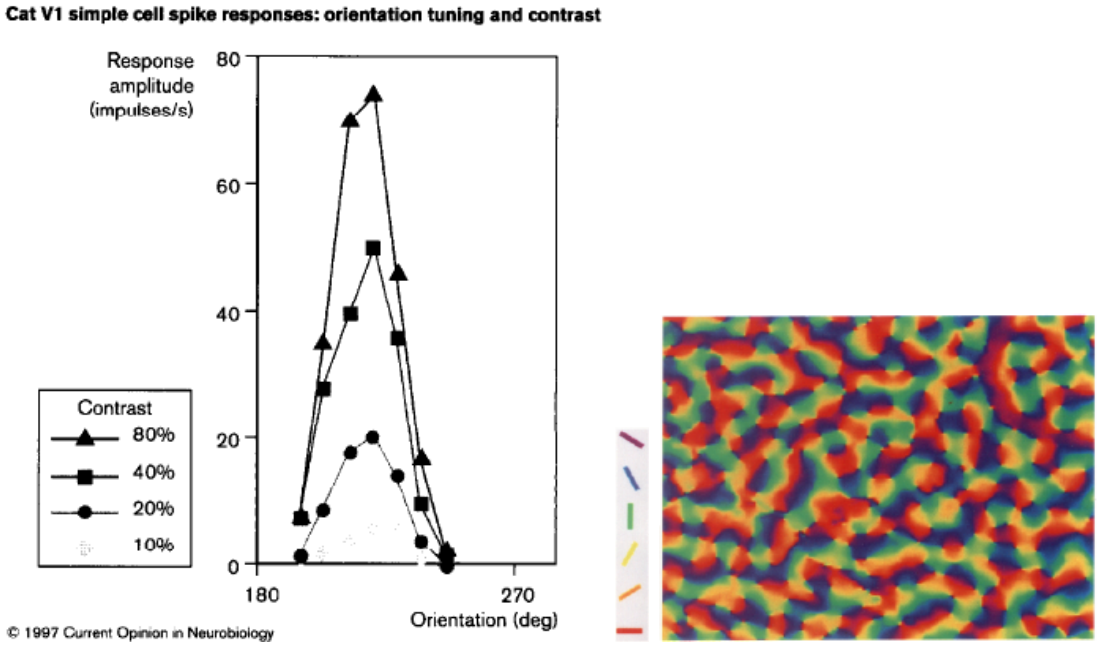

Figure 2.3: **Left**: *orientation tuning in cat V1. The mean response of a cell to a full range of orientation is plotted for different constrasts, showing that the tuning width is constrat invariant. Taken from [28].* **Right**: *an orientation map in adult macaque monkey V1, taken from [5]. The map spans a 7.5 mm $\times$ 5.5 mm area on the surface of the cortex. Each neuron is colored according to its preferred orientation, using the color key on the left. Preference gradually changes, forming a smooth map composed of "iso-orientation patches". The typical distance between two such blobs of same orientation preference is 1 mm.*

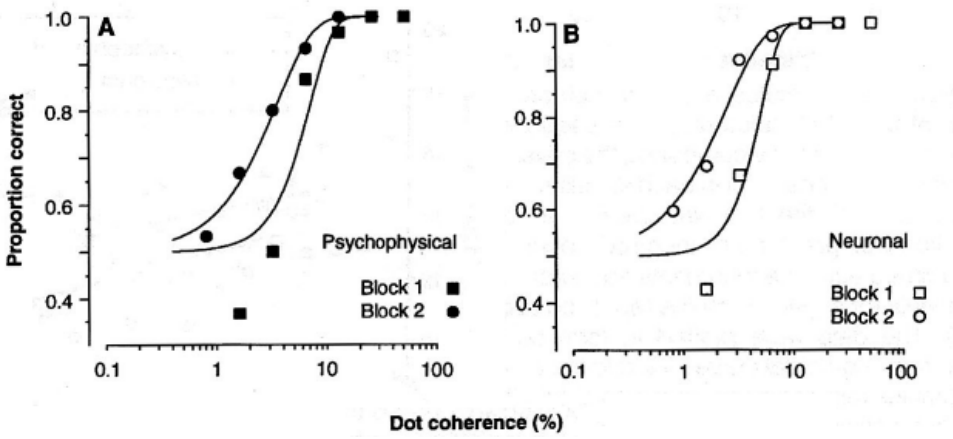

Figure 2.4: *The striking correspondance between psychophysical and neuronal level of performance, as well as patterns of improvement, in a coarse discrimination task involing the motion direction of a random dot field. The x-axis can be understood as the level of difficulty of the task (the task gets easier and easier as the dot coherence increases). Reprinted from [69].*

sized performance of an ideal observer based on the responses of single neurons. The neuronal sensitivity mirrored the perceptual sensitivity with a rather striking precision (see figure 2.4).

More generally, a wide range of single neurons experiments have been carried out in W. T. Newsome's laboratory. One of their purposes, as stated in [45], is to understand the link between the activity of single cells and perception. It is a first step, in their view, to understand how information can be combined by many cells. Their experimental data, such as that of [69] presented above, show that "significant number of [single] neurons perform at levels that compare favorably with the overall behavioral performance of the organism".

No changes in receptive field properties

Contrary to what happens in audition and somatosensation, no major changes in the basic properties of the neurons' receptive fields (e.g. location, size, orientation) have been reported in the visual cortex. R. E. Crist and coworkers found that training monkeys in a bisection discrimination task did result in a dramatic perceptual improvement, while receptive fields in V1 neurons remained unchanged. In addition, "visual topography was indistinguishable between trained

and untrained animals" [10].

Therefore, neural changes, if any, have to be found in more subtle response properties such as the tuning amplitude or width.

Gain modulation

The first physiological investigations in the neural basis of improvements in orientation discrimination have shown that specific gain modulation occur in V1 and V2. A gain depression for neurons tuned at and around the trained orientation was reported in abstract form by G. M. Ghose [20]. It was confirmed a few months later by Schoups and collegues [55]. In contrast, no gain modulation was found in [56] and [21], neither V1 nor in V2. Furthermore, a gain amplification was found in V4 [68].

The contradictions between these electrophysiological recordings is also mirrored in the domain of fMRI, where a couple of studies have shown both gain amplification and depression in V1 (see for example [17]).

Local sharpening in V1

Further studies have focussed on the sharpness of the tuning. A. Schoups found that, in V1, neurons tuned at and around the TO exhibited an increase in the slope of their tuning curves at the TO (see figure 4.3, page 34). This may stem from a sharpening of the tuning curves, although the widths of the tuning curves after learning are not reported.

These results were challenged by another study by G. M. Ghose [21] in which "no conclusive changes that could account for learning at the behavioral level had been demonstrated".

Local sharpening in V4

In V4, T. Yang and J. Maunsell reported a global sharpening of the tuning curves, with a more prominent effect for those neurons tuned at and around the TO. Similarly, S. Raiguel and collegues found similar patterns of changes as A. Schoups found in V1 (see above) but more pronounced. They concluded that

learning mainly modifies the tuning properties of the most informative neurons in V4 [49]. Moreover, a complementary study suggested that this increase in slope is more likely to be the result of localized and asymmetric deformations of the tuning curves on the side of the TO than of a symmetric sharpening. This was done by comparing the slope of the tuning curves at the trained orientation with that of the symmetrically opposite point on the other side of the preferred orientation, which proved to be different for neurons tuned in a range of 22 to 67 degrees away from the TO.

Chapter 3

Population encoding

3.1 Introduction

In the previous chapter (see page 12), we have described the general physiological properties of the neurons in the visual cortex. In the following few chapters, we need a model of these neurons that should provide a relevant basis for further computational explorations in perceptual learning.

As we have seen, in the domain of direction of motion discrimination, several studies have shown that under certain conditions, only one neuron can be responsible for the behavioral performance of the subject. However, we know that information is encoded in large populations of neurons. Hence the use of a *population coding* approach. It is a simple model where we do not worry about the underlying connectivity (unlike [65, 57]), but in which all the important elements of the encoding problem are included: the representation is indeed distributed, noisy, does possibly include correlations, and the response properties of each neuron directly reflect what is observed in neurophysiology (tuning curves). This model can in fact be looked on as the simplest reasonable model of encoding that allows the use of a decoding strategy on top of it. *Population decoding* will be introduced in chapter 5 and used in chapter 6 in order to relate perception to behavior.

3.2 General setting

Within the framework of a population code, one (or more) feature of the visual stimulus (here, orientation) is encoded by N neurons, through their joint activities, taken to be a vector of firing rates.

Each cell is characterized by its *mean response* to a given stimulus as well as by the *variability* of its response. Therefore, cell number i is modelled by two components:

- its tuning curve, which is a bell-shaped function giving the mean firing rate $f_i(\theta)$ – number of spikes fired in a given time window – in response to a certain orientation θ
- the probability density function $p_i(r_i|\theta)$ of evoking a firing rate r_i given a stimulus θ. It incorporates the stochastic aspect of the response. The activity of the cell is thus a random variable with a dynamic mean $f_i(\theta)$.

A common and often fair approximation is that of noise independence between neurons. In this case, the probability of a stimulus with orientation θ evoking the vector of responses $\mathbf{r}$ can be written as

$$p[\mathbf{r}|\theta] = \prod_{i=1}^{N} p_i(r_i|\theta) \tag{3.1}$$

3.3 Tuning curves

According to physiological recordings, the function $f_i(\theta)$, called tuning curve, is often well approximated by a *Gaussian*. In which case we have this four-parameters model:

$$f_i(\theta) = b_i + f_i^{\max} \exp\left[-\frac{(\theta - \theta_i^{\text{pref}})^2}{2\sigma_i^2}\right] \tag{3.2}$$

- θ_i^{pref} is the preferred orientation of the neuron, i.e. the stimulus that evokes the maximum response. The distribution of θ_i^{pref} (in other words, the orientation spectrum of the orientation map in V1) has been shown not to be flat: more neurons code for the horizontal and vertical orientations (the so-called "principal orientations") than for others. In our simulations, nonetheless, we often assume that preferred orientations are evenly distributed.

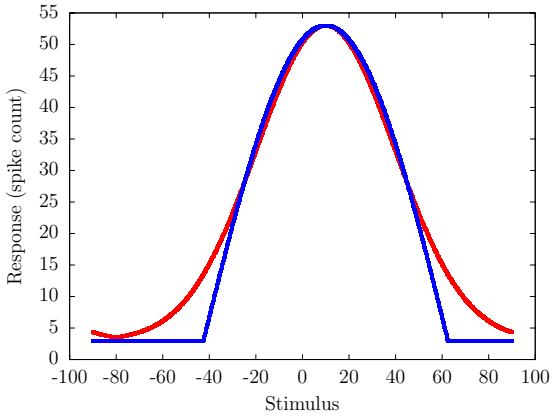

Figure 3.1: *Two kinds of tuning curves. Red: Gaussian (eq. 3.2), with* $\theta^{pref} = 10$ *deg,* $b = 10$ *spikes,* $f^{max} = 50$ *spikes, and a width at half-height equal to* 70 *deg. Blue: rectified cosine (eq. 3.3), with same parameters.*

- the baseline b_i is the average firing rate in response to non-optimal stimuli. Typically, it is around $10 - 15$ spikes per second, in monkey V1 ([68]).
- $f_i^{\max}$ is the maximum firing rate, such that at $\theta = \theta_i^{\text{pref}}$, the neuron fires on average at $b_i + f_i^{\max}$. In monkey V1, $f_i^{\max}$ is about 20 spikes per second.
- σ_i characterizes the width of the tuning curve. It is related to the width at half-height of the tuning curve without baseline, W_i, by the following equation:

$$W_i = 2\sigma_i\sqrt{2\log 2}$$

It is also possible to use *rectified cosine* tuning curves:

$$f_i(\theta) = \begin{cases} b_i + f_i^{\max}\cos\left[\frac{2\pi}{3W_i}(\theta - \theta_i^{\text{pref}})\right] & \text{if } |\theta - \theta_i^{\text{pref}}| < \frac{3W_i}{4} \\ b_i & \text{otherwise} \end{cases} \tag{3.3}$$

Figure 3.1 (left) provides an example for both types of tuning curves.

3.4 Neural variability

When physiologists record the firing rate of cell i in response to a stimulus θ, the observed value differs from trial to trial. The brain is indeed inherently noisy. When recording the tuning curves, physiologists actually average over trials, so that $f_i(\theta)$ is indeed the averaged firing rate in response to orientation θ. But in

our framework, we need to *simulate trials*. To capture this stochastic aspect of the neuronal responses, we generate each trial by sampling from a distribution centered in $f_i(\theta)$. This distribution can be a Poisson distribution, for instance. In which case :

$$p_i(r_i|\theta) = \frac{f_i(\theta)^{r_i}}{r_i!} \exp(-f_i(\theta)) \tag{3.4}$$

Note that r_i is taken to be a spike *count*, not rate. If one wants to work with firing rates, r_i is simply obtained by multiplying the rate by the temporal window Δt. In this dissertation, we will only consider spike counts.

Poisson statistics describes pretty well the noise in neurons with low spike counts in the given time window. For higher activities, the noise is better represented by a Gaussian distribution:

$$p_i(r_i|\theta) = \frac{1}{\sigma_n\sqrt{2\pi}} \exp\left[-\frac{(r_i - f_i(\theta))^2}{2\sigma_n^2}\right] \tag{3.5}$$

Poisson statistics has a Fano factor (variance to mean ratio) equal to unit, that is, the variance equals the mean. According to neurophysiological recordings, the Fano factor is more likely to lie between 1 and 2, which can be modelled in the case of Gaussian noise by setting the variance σ_n^2 different from the mean $f_i(\theta)$.

An example of a trial in a population of 100 neurons is depicted in figure 3.2.

3.5 Code accuracy

Whenever scientists work with codes, they need to have some way to assess their accuracy. For example, in digital information transmission, the quality of the code is given by the mutual information (also called Shannon information) between the source and the receiver. This information is "input independent" (the variability due to the variability in the input is indeed averaged away), and as such, is not well-suited to population codes. When working with neural population codes, one would like to know, *given a stimulus* and a subsequent response from the population, what certainty we have about the estimate of the stimulus from the response. Chapter 5 will present various ways of decoding the population response (i.e. guessing the underlying stimulus), but here we want to know, independently of the decoder, how much information about the stimulus is carried by the code.

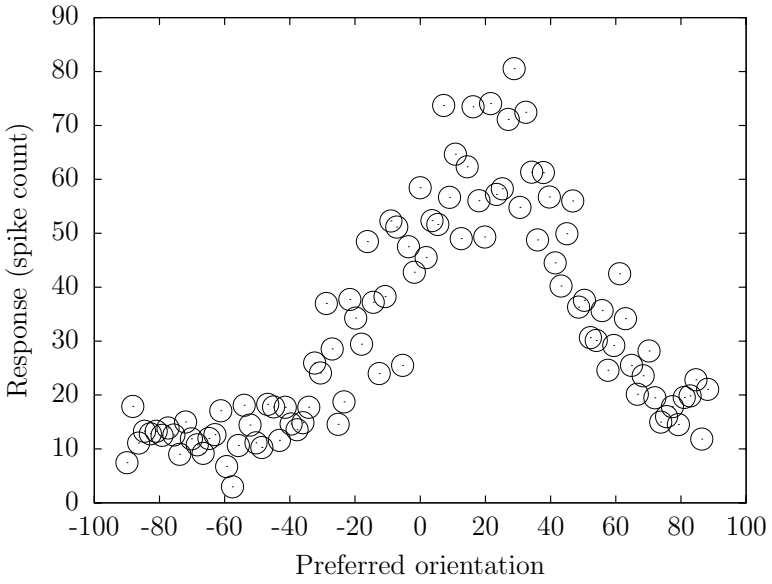

Figure 3.2: *One "trial", i.e. the activity of the population – generated by the model – in response to a stimulus (here $\theta = 20$ deg). Firing rates of 100 neurons are sampled from Gaussian distributions (eq. 3.5) with a Fano factor equal to 1.3. Tuning curves are Gaussians with same parameters as in figure 3.1.*

An appropriate measure of this information is *Fisher information*. It is defined as the expected value of the curvature of the log-likelihood of the response given the stimulus:

$$I_F(\theta) = -\int p(\mathbf{r}|\theta)\frac{\partial^2 \log\left(p(\mathbf{r}|\theta)\right)}{\partial\theta^2}d\mathbf{r} \tag{3.6}$$

Simpler expressions can be derived when we know the analytical form of the log-likelihood, namely in the case of Poisson or Gaussian variability:

$$\textbf{Poisson noise} \qquad I_F(\theta) = \sum_{i=1}^{N}\frac{f_i'^2(\theta)}{f_i(\theta)} \tag{3.7}$$

$$\textbf{Gaussian noise (Fano k)} \qquad I_F(\theta) = \frac{1}{k}\sum_{i=1}^{N}\frac{f_i'^2(\theta)}{f_i(\theta)} + \frac{1}{2}\sum_{i=1}^{N}\frac{f_i'^2(\theta)}{f_i^2(\theta)} \tag{3.8}$$

The latter equation actually holds assuming that the variance of the neural response depends linearly on its mean, for each neuron. This is a commonly used simplifying assumption, hence the use of the Fano factor as a noise characterization (variance = mean $\times$ Fano factor).

A very important theorem related to Fisher information is the Cramer-Rao theorem. It states that

> **The variance of any unbiased estimator of the real value θ has variance greater than the inverse Fisher information at θ:**
>
> $$\text{if} \quad \left\langle \hat{\theta} \right\rangle = \theta \quad \text{then} \quad \sigma^2(\theta) \geq \frac{1}{I_F(\theta)} \tag{3.9}$$

We will see that some estimators may sometimes have a lower variance than that allowed by the bound, in which case they are actually strongly biased. As we shall see later as well, under the assumption that the stimulus is indeed decoded somewhere at some point in the brain, the psychophysical performance in a typical discrimination task only depends on the bias and variance of the estimator. Therefore, provided the estimator is unbiased, Fisher information may be looked on as a measure of discrimination performance.

In fact, a more general version of the theorem may be stated as follows ([8]):

> **The variance of any estimator of the real value θ has a bias b and variance σ^2 that are such that:**
>
> $$\sigma^2(\theta) \geq \frac{(1 + b'(\theta))^2}{I_F(\theta)} \tag{3.10}$$
>
> where b' stands for the first order derivatives of the bias with respect to the stimulus.

Chapter 4

A mechanistic model of neuronal fine re-tuning

4.1 Introduction

We here want to come up with a *mechanistic model of perceptual learning*. More precisely, our model should meet the following three requirements:

1. the encoder must be learnt in a way that is highly dependent on the input distribution. Indeed, we want to further consider the idea that the distribution of stimulus in the task – usually strongly biased toward the trained orientation – plays an important role in learning

2. moreover, we want the learning algorithm to be *unsupervised*: no feedback must be needed to learn the new representation. We indeed want to stay consistent with the psychophysical finding that feedback is actually not necessary for perceptual learning to occur

3. the re-representation of the stimulus must have the same properties as the former representation. Particularly, neurons involved in the "output" of the learning process must exhibit orientation selective and noisy responses, for reasons that will become clear later (we actually want to be able to apply the same decoding schemes – as described in chapter 5 – to evaluate the new representation in terms of psychophysical performance).

In this chapter, we first re-express the above criteria in more practical terms, and we present a model that meets all of them. As we shall see, nonetheless, our first attempt is not successful in generating the same patterns of neural fine re-tuning as recorded by physiologists. We then explain how the first model must be incremented in order to eventually predict the tuning curves modifications. We present the experiments and results, before discussing further issues.

4.2 Our first attempt: a self-organizing string of neurons

The architecture of the first model we use is depicted in figure 4.1. We know that in the primary visual cortex, the encoding of the stimulus angle is very well modelled by tuning curves and variability functions (chapter 3). We want to incorporate this into our model. Therefore, we have a first layer equivalent to the populations we have seen so far: a bench of neurons with preferred orientations evenly distributed between -90 and 90 degrees. The input of the system can thus be a mere floating point number representing the angle (stimulus). The resulting activity of the first layer is what we called a "trial" (e.g. see figure 3.2, page 23). Due to the neural variability, this response is always noisy.

Our model is supposed to learn the encoding of the stimulus other than this primary encoding in V1. We thus need a second layer that will learn a new representation of the stimulus from the first layer. We now explain our choice for this architecture and for the learning algorithm by reformulating the criteria expressed in the introduction.

4.2.1 Reformulating the criteria

In order to meet the third criterion as stated in the introduction, the neurons in the second layer should be orientation selective and noisy as well. This naturally leads us to use what is called a "self-organizing" algorithm. The second layer auto-organizes so that, eventually, each of its neurons becomes strongly activated by a continuous subset of neurons in the first layer. Consequently, since the first layer is continuously orientation selective, neurons in the second layer will

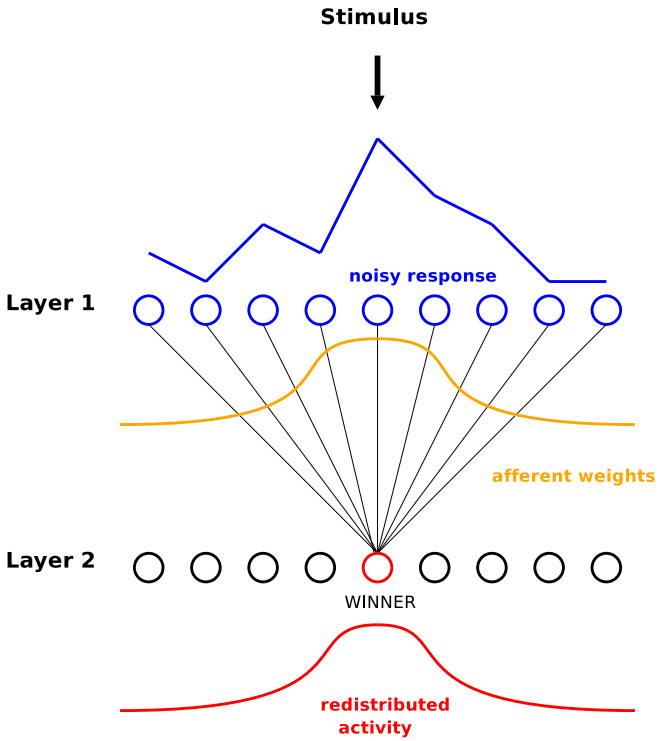

Figure 4.1: *Architecture of the self-organizing string of neuron described in the text.*

be as well. Finally, the noise that comes from the first layer will naturally be propagated forward to the second layer[1]. For concreteness, we assume that the second layer is fully interconnected with the first one. We also assume, without loss of generality, that the numbers of neurons in the first and second layers are identical (N).

In order to meet the first criterion, the new representation learnt by the second layer must differ from one prior distribution of inputs to another. The second layer should therefore learn from the statistics of the first layer, which directly reflect the statistics of the input stimulus of course. This naturally conducts us to use Hebbian learning. To quote [11]:

> *In 1949, Donald Hebb conjectured that if input from neuron A often contributes to the firing of neuron B, then the synapse from A to B should be strengthened.*

Another way to express this principle is:

> *[...] synapses change in proportion to the correlation or covariance*

[1] although quite attenuated because, as we shall see, the second layer acts like a filter of the first layer.

of the activities of the pre- and post-synaptic neurons.

Different learning rules have been designed that follow this general principle. The learning rule we use here is added some competition between neurons, which has both benefits of making the learning process more stable and allowing self-organization of the second layer.

Finally, the second criterion is achieved by using an unsupervised learning rule.

4.2.2 The Self-Organizing Map

A very well suited algorithm that meets all the above is Kohonen's Self-Organizing Map ([35]). This algorithm[2] is originally described (and more likely to be known) in the case of two-dimensional inputs. Here we adjust it for the needs of *one-dimensional strings of neurons.*

In this model, the second layer is fully connected with the first layer, and we denote by $\mathbf{w_a}$ the matrix of synaptic weights. When a stimulus is presented, θ, the first layer has activity $\mathbf{r}(\theta)$, and the subsequent activation of the second layer $\mathbf{a}(\theta)$ is merely the weighted sum of $\mathbf{r}(\theta)$ using the synaptic weights:

$$\mathbf{a}(\theta) = \mathbf{w_a} \cdot \mathbf{r}(\theta) \tag{4.1}$$

After this activity is computed, a "winner" is picked up among neurons in the second layer. The winner is that neuron with highest activity. The activity in the second layer is subsequently *redistributed* around the winner, following a Gaussian profile. Neurons close to the winner are thus given a strong activity, although they were not necessarily much activated by the propagation of the input signal only. We denote by $\mathbf{a}'(\theta)$ this redistributed activity:

$$a'_i(\theta) = \exp\left[-\frac{(i - i_w(\theta))^2}{2\sigma_a^2}\right] \tag{4.2}$$

where $i_w(\theta)$ is the index of the winner. The difference between i and $i_w(\theta)$ should be understood as a circular distance, since the network is toroidal. Finally, the weights of the afferent connections are updated according to a normalized Hebbian

[2] we should actually say "this family of algorithms", for the original SOM developed into many versions. We here suggest one of them.

learning rule:

$$w_{ij} = \frac{w_{ij} + \varepsilon \cdot r_j \cdot a'_i}{\sqrt{\sum_{k=1}^{N} (w_{ik} + \varepsilon \cdot r_k \cdot a'_i)^2}} \tag{4.3}$$

where ε is a small learning rate. Each line of the weight matrix has unit Euclidean norm.

The Hebbian learning rule tends to push the weight vector of the most active neuron and its neighbors toward the input vector (by maximizing the dot product). In other words, neurons become selective and get a "preferred stimulus", learn to better represent it, and as a result of the bell-shaped redistribution of activity, neurons close to each other end up having similar preferred stimuli. Therefore, this algorithm is suitable for the production of a topographically self-organized string of neurons selective to orientation.

For computational efficiency, we normalize the response of the first layer (so that they have unit Euclidean norm) before computing the activity in the second layer. We also normalize in the same way the subsequent redistribution of activity in the second layer, before applying the learning rule.

For the sake of simplicity, we bias the initial random weight matrix by reinforcing its diagonal. Indeed, if the self-organizing algorithm ensures that the first layer is topographically mapped down to the second layer (neighboring neurons in the string have neighboring preferred orientations in the input space), neuron i in the first layer will not necessarily correspond to neuron i in the second layer. Up to a circular shift of the weight matrix, this can be the case. Therefore, without loss of generality, we can force the system into developing a "human-readable" weight matrix.

4.2.3 Computing the resulting tuning properties

Once the network has properly self-organized, the neurons in the second layer become orientation-selective. Each neuron is therefore characterized by its tuning curve and the variability of its responses around the tuning curve.

The tuning curve of neuron i in the second layer, evaluated at stimulus θ, is denoted by $f_i(\theta)$. It is computed by presenting stimulus θ to the first layer many times (3000 here). Each time, the activity is propagated to the second layer to get

$a_i(\theta)$, and $f_i(\theta)$ is obtained by averaging over trials. The variability of neuron i is computed similarly. We think, although not show, that if the noise is Gaussian in the first layer, it is also Gaussian in the second layer (the activity in the first layer is a Gaussian vector, and any linear transformation of a Gaussian vector is a Gaussian vector as well). Hence the only need of computing the variance of the response around the mean (over trials, generated in the same way as mentioned above).

4.2.4 Simulation

We now incorporate this learning algorithm into the architecture in figure 4.1. We run the self-organizing process using a flat distribution of orientations, and compute the tuning curves (termed "tuning curves before learning") and the variabilities as described above. Then, taking the network in its final state (no resetting – which is supposed to correspond to the stable state of a naive population), we run the network using a prior probability of inputs that favors the trained orientation. We then compute the tuning curves again ("tuning curves after learning"). The prior is Gaussian with spread σ_p. We cannot give any formula, since the process of sampling a stimulus is a bit more complicated than a single Gaussian sampling because of the circularity of the variable. What actually really matters is to know that the input distribution is biased toward the trained orientation, and that σ_p is a measure of how strongly biased it is.

The parameters of the first layer population are given in the table below.

Parameter	**Value**
N	100
Baseline	0
Amplitude	70 spikes
Fano factor	1.3
Trained orientation	0 deg.
$W_i^{(0)}$	40 deg. for all i
σ_r	40 deg., decreased to 5 deg.
σ_p	40 deg.
η	0.05, decreased to 0.01

The learning rate is large at the beginning of the training (0.05), whereas it is manually lowered to 0.01 after 5,000 iterations, so as to make the topographical organization of the weights more precise. Similarly, the activity profile in the second layer is also made broader in the beginning ($\sigma_r = 40$) and sharper in the end. Every 10 iterations it is multiplied by 0.99 and saturates to 5. Stable patterns of weights are obtained after 10,000 iterations. Simulating perceptual learning – i.e. running the SOM algorithm with a peaky prior probability of input stimulus – takes a bit more iterations (50,000).

4.2.5 Results

The results of the self-organizing process are shown in figure 4.2, for both flat and peaky priors. The density of neurons tuned close to the TO increases, which can be seen from the distortion of the weight matrix (top plot): the centers of the bell-shaped weights profiles are attracted toward the TO.

Direct "geometrical" consequences of this neural recruitment are a sharpening of the bell-shaped afferent weights for those neurons tuned around the TO and a widening for neurons tuned farther from the TO[3]. The sharpening and broadening propagate to the resulting tuning curves of course (see figure 4.2, bottom left). Tuning curves are indeed a bit sharper after than before learning, for neurons tuned at and around the TO, whereas they are much broader for other neurons.

Since the weights are normalized so that they have unit Euclidean norm, the amplitude of these "weight blobs" drops as they get wider, and vise-versa. Subsequently, in terms of activity, there is a gain amplification for central neurons and a gain depression for peripheral ones (see the tuning curves after learning depicted in figure 4.2, middle left). This result is in contradiction with physiological data, as reported in abstract form by A. Schoups in [55]. She found a decrease of activity for those neurons.

As to the neural variability, the main effect of the peaky prior distribution of input is a significant increase in the Fano factor for neurons tuned at and around the TO, and a similarly important decrease for neurons tuned around the orthogonal orientation (figure 4.2, bottom right).

[3]the reader may want to compare the horizontal blue area in the top or bottom rows of the right hand side matrix in figure 4.2 with that of the central rows.

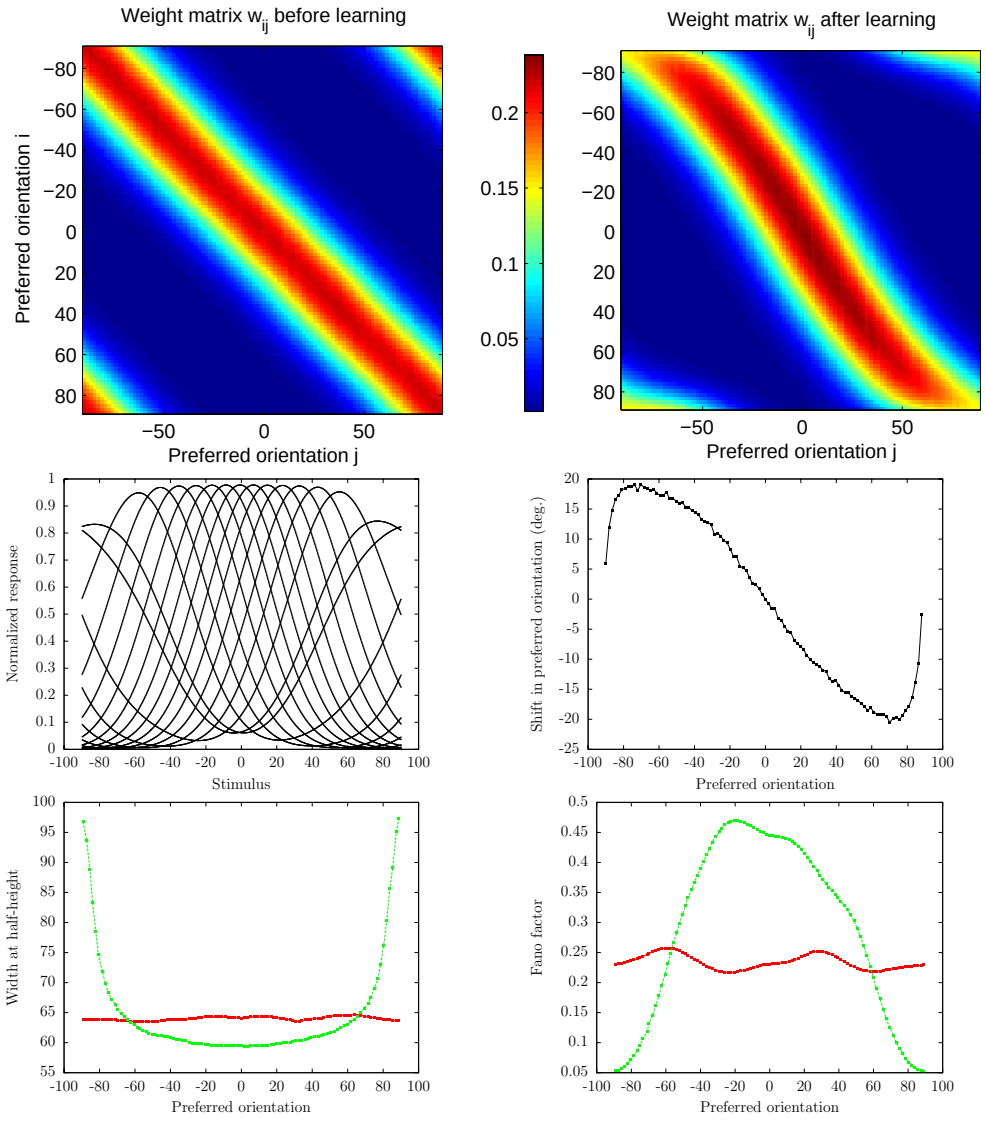

Figure 4.2: **Top left**: *the afferent connections from the first layer to the second layer, when the SOM algorithm has converged (after 10,000 iterations) with a flat distribution of input stimuli.* **Top right**: *the same connections after another 50,000 iterations with a peaky prior probability of input centered at the TO (0 degrees here). The two plots share the same scale.* **Middle left**: *tuning curves after learning. Tuning curves are "normalized" in the sense that the weights were learnt from normalized inputs and normalized redistribution of activity in the SOM algorithm, in addition of being normalized themselves.* **Middle right**: *the shift in preferred orientation for each neuron (identified by its preferred orientation on the x-axis).* **Bottom left**: *the width of the tuning curve of each neuron, before learning (red) and after learning (green).* **Bottom right**: *the Fano factor of each neuron. It is computed by averaging* $\sigma_i^2(\theta)/f_i(\theta)$ *over* θ*, where* σ_i *is the variance of the response computed as described in the text. In addition, we have multiplied the previous quantity by the constant* f^{max} *(amplitude of the tuning curves of the neurons in the network's first layer). It is therefore possible to also compare the Fano factors between layers (*$Ff = 1.3$ *in the first layer).*

4.2.6 Comparison to physiological data

To allow a direct comparison between the new tuning curves and that reported by [56] following intensive training in monkey (see page 16), we plot the slope of the tuning curves at the TO for each neuron (figure 4.3, top left), as A. Schoups did. The slope increases for all neurons tuned farther than 30 deg away from the TO. However, contrary to the observations in [56], the slope drops after learning for those neurons that used to have the highest slope ("most informative neurons"). Similarly, we plot the signal-to-noise ratio at the trained orientation, for each neuron (figure 4.3, bottom left). We observe that this quantity increases after learning. The difference is larger for neurons tuned close to the TO. On the contrary, Schoups and coworkers found a significant *decrease*.

4.2.7 Discussion

Qualitatively, the shift in orientation preference is similar to that observed in the auditory and somatosensory cortices (as we saw in chapter 2, page 11). The main learning effect in our model is indeed the enlargement of the area of cortex that responds to the trained stimulus. This cortical reorganization, however, cannot really be interpreted quantitatively. Indeed, sharpening or broadening the prior distribution of input has been observed to induce a direct augmentation or diminution of the shift, and the spread of the prior we use is somehow subjective. This comment holds for all other changes reported above: the sharpness of the prior controls the intensity of the learning effects.

We can first conclude that a self-organizing string of neurons driven by unsupervised Hebbian learning can encode the prior probability of the input stimulus through the distribution of its neurons' stimulus preferences, thus replicating the neural correlates of perceptual learning in audition and somatosensation modalities [50, 51]. It fails, however, to fully predict the neural changes observed in the early visual areas following intensive practice. Tuning curves do not sharpen nor broaden in the same fashion as was reported by A. Schoups and her colleagues, and qualitatively confirmed by further studies [68, 49]. Hypothesized gain modulations for neurons tuned near the TO [20, 55] are not replicated neither. Finally, the model comes up with an untested and presumably not easily

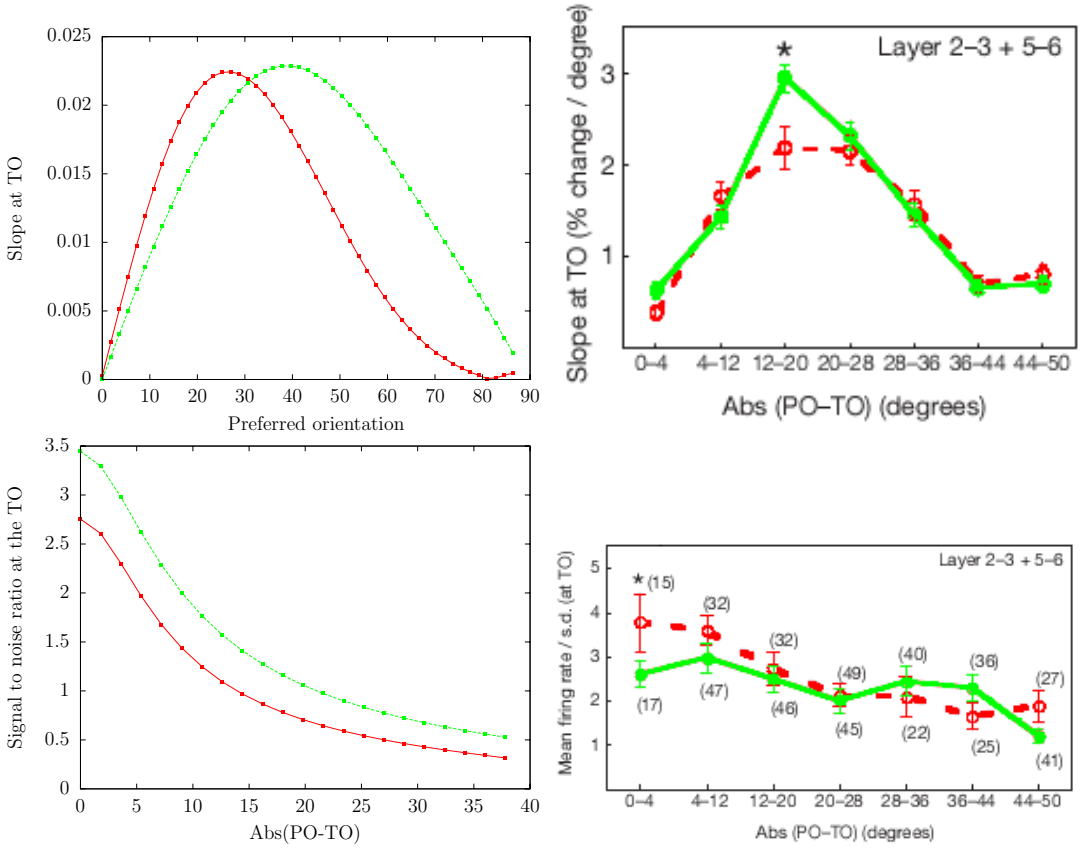

Figure 4.3: **Top left**: *the slope at the trained orientation of the tuning curve of each neuron, in our simulation, before learning (red) and after learning (green). Each neuron is represented by the distance of its preferred orientation to the TO (Abs(PO-TO)).* **Bottom left**: *the signal-to-noise ratio of each neuron* at the trained orientation. **Top and bottom right**: *the same data obtained by Schoups* et al. *from physiological recordings. Reprinted (and recolored) from [56].*

testable physiological prediction: a shift of orientation preference (cf page 15). An optimistic reader would acknowledge, however, that the model used herein does predict *specific patterns of neuronal fine re-tuning*, although these are in partial contradiction with physiological data.

The learning rule we have used in this section meets the requirements we set up in the introduction. Still, it has a number of shortcomings and biological incompatibilities. Picking one winner in the whole population, for example, is artificial. So is the redistribution of activity according to the Gaussian profile. It has been suggested that perceptual improvement would partially stem from a modification of the lateral interactions among neurons involved in the task [9, 23, 24], as an alternative to – or in addition of – other top-down influences. To quote [9]:

> [The degree of specificity shown by perceptual learning] *suggests that the learning cannot be achieved by cortical recruitment alone, as proposed in current models, but is likely to involve a refinement of lateral interactions within the cortex and possibly a gating of lower level changes by attentional mechanisms.*

In SOM, the redistribution of activity is somehow equivalent to fixed lateral interactions. In the next section, we investigate further the idea of learning lateral connections.

4.3 Learning lateral connections

Our first attempt failed to fully replicate neurophysiological findings. In particular, sharpening of the tuning curves was relatively modest. The main reason why SOM does not allow much sharpening of the tuning curves may reside in the static aspect of the lateral connections. We here want to augment the previous model with dynamic lateral interactions. To do so, we get inspired by the LISSOM model, described below.

4.3.1 LISSOM

LISSOM ([42, 43, 63]) stands for "Laterally Interconnected Synergetically Self-Organizing Map". It takes after the original SOM model described above, and

incorporates more biologically plausible lateral interactions among neurons in the map. These lateral interactions are also able to learn. LISSOM has proved very good at modeling some aspects of the mammalian visual cortex, e.g. ocular dominance, neuronal selectivity to orientation, to direction of motion, in a way that reflects the topographic mapping of the retina onto the cortex.

Warning Shrinking SOM to a one-dimensional model with the same learning algorithm, did not represent, in our opinion, a big change in the concept embodied in SOM. SOM is indeed usually looked on as a "machine learning mechanism", where a population of units learns to represent input patterns as accurately as possible. Therefore, it does not really matter whether the data is one- or two-dimensional. In contrast, LISSOM was primarily intended to model small pieces of visual cortex, that is, sheets of neurons. Reducing LISSOM to a one-dimensional network thus sounds a bit more ambiguous, for 2D-retinotopy is one of the key features of the model. Instead of saying we are using LISSOM, we would rather say we are using LISSOM's algorithm while drastically changing its architecture.

Again, this reduction of the LISSOM model is justified by:

- the need of having as input a stimulus *value* θ, as opposed to a two-dimensional retinal pattern of activation oriented at θ
- the need of having as output a population code of the stimulus value, so as to apply decoding strategies to assess perceptual performance as described in chapter 5.

Although these requirements could be fulfilled by using the traditional 2D LISSOM architecture, but it would obviously not be the easiest way to meet them. It will be for future studies to work on a more realistic two-dimensional model where LISSOM's most powerful features would be fully exploited.

Architecture The overall mechanism is depicted in figure 4.4. In LISSOM, the afferent connections are different from that of SOM: each neuron in the second layer receives input only from a continuous subset of neurons in the first layer. Neurons therefore have "receptive fields". Furthermore, each neuron in the second layer is connected to a continuous subset of neurons in the same layer. More

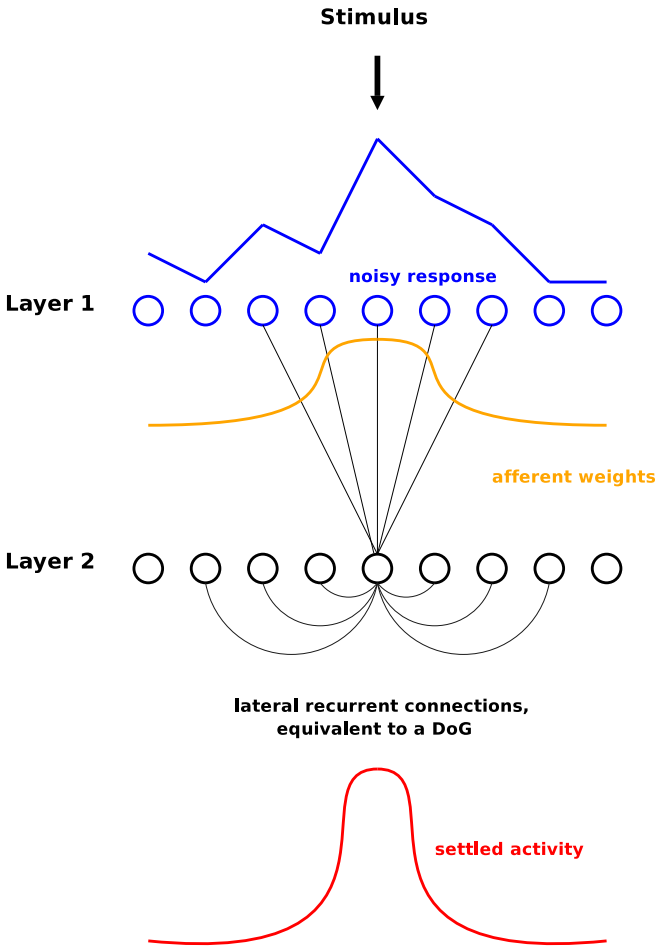

Figure 4.4: *LISSOM's algorithm applied to one dimensional populations of neurons.*

precisely, considering one neuron, nearby neurons send excitatory input to it, whereas neurons within a wider range send inhibitory signal. All these connections can be summarized in three weight matrices: $\mathbf{w_a}$ for the afferent, $\mathbf{w_e}$ for the excitatory, and $\mathbf{w_i}$ for the inhibitory connections. Each matrix is in fact null except within a certain radius (α_a, α_e or α_i) from its diagonal (which corresponds to the receptive fields). The use of matrices simplifies the notations in what follows.

Response generation Activity in the second layer is computed in two steps:

- first, the activity $\mathbf{r}$ in the first layer propagates to the second layer through the afferent connections (like in SOM). The afferent activity is thus given by

$$\mathbf{a}(\theta) = \gamma_a(\mathbf{w_a} \cdot \mathbf{r}(\theta)) \tag{4.4}$$

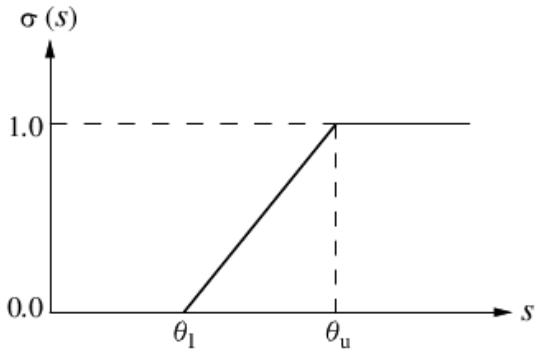

Figure 4.5: *The activation function we use. θ_l and θ_u are to be adjusted so that the activity of each neuron stays in the linear part of the function (no saturation). Reprinted from [44].*

where γ_a is a constant scaling factor. The afferent weights matrix is now denoted by $\mathbf{w_a}$.

- this afferent contribution is further processed by lateral connections (recurrent process), and the activity in the second layer settles down according to the following equations:

$$\begin{aligned} \text{First iteration} \quad & \mathbf{a'}_0 = \sigma\left[\mathbf{a}\right] \\ \text{Iteration } i \quad & \mathbf{a'}_i = \sigma\left[\mathbf{a} + \gamma_e(\mathbf{w_e} \cdot \mathbf{a'}_{i-1}) - \gamma_i(\mathbf{w_i} \cdot \mathbf{a'}_{i-1})\right] \end{aligned} \tag{4.5}$$

  A few iterations (we used 5) are sufficient for the activity to reach a stable state. σ is a piecewise linear approximation of a sigmoid activation function that keeps the activity between 0 and 1 (see figure 4.5). γ_e and γ_i are constant scaling factors that determine the relative strength of excitatory and inhibitory connections. The fact that inhibitory connections extend to a larger radius compared to excitatory connections is crucial to ensure the stability of this settling process.

Learning the afferent connections The afferent connection from neuron j to neuron i is learnt according to the same normalized Hebbian learning rule as in SOM:

$$w_{a_{ij}} = \frac{w_{a_{ij}} + \varepsilon_a \cdot r_j \cdot a'_i}{\sqrt{\sum_{k=1}^{N} \left(w_{a_{ik}} + \varepsilon_a \cdot r_k \cdot a'_i\right)^2}} \tag{4.6}$$

Learning the lateral connections The lateral connections are allowed to learn in the same way:

$$w_{e_{kl}} = \frac{w_{e_{kl}} + \varepsilon_e \cdot \eta_k \cdot \eta_l}{\sqrt{\sum_{k=1}^{N} \left(w_{e_{kl}} + \varepsilon_e \cdot \eta_k \cdot \eta_l\right)^2}} \tag{4.7}$$

and

$$w_{i_{kl}} = \frac{w_{i_{kl}} + \varepsilon_i \cdot \eta_k \cdot \eta_l}{\sqrt{\sum_{k=1}^{N} \left(w_{i_{kl}} + \varepsilon_i \cdot \eta_k \cdot \eta_l\right)^2}} \tag{4.8}$$

γ_a, γ_e and γ_i are the learning rates for afferent, excitatory and inhibitory connections respectively.

Adapting the parameters of the activation function θ_l and θ_u (see figure 4.5) must be updated regularly so that the neurons do not saturate. Usually, it is done manually. Other mechanisms can achieve it in a more biologically plausible and automatic way, such as homeostatic adaptation of the level of excitability of each neuron [66]. During the course of the project, we investigated both possibilities. For now, we present the results we obtained by systematically rescaling θ_l and θ_u – at every learning iteration – so that the *afferent* activity stays in the good range. Although it may not seem a very realistic option, it brought us the most interesting results. At the end of this chapter, we discuss the implications of using homeostatic adaptation.

4.3.2 Further simplifications

In the original LISSOM model, all the connections have a limited spatial extent (receptive fields). Here, we choose not to restrain the receptive fields, and we explain this choice in the discussion at the end of the chapter. Therefore, all neurons are fully interconnected, via both afferent and lateral connections. However, in doing so, we face another problem. Lateral excitatory and inhibitory connections learn following the same learning rule. Therefore, if they have the same spatial extent, they will end up being completely identical. The combined interactions ($\gamma_e \mathbf{w_e} - \gamma_i \mathbf{w_i}$) will therefore not be a good "Mexican-hat" profile, required for proper and stable self-organization. This is addressed by keeping the excitatory

connections constant, and extending to a smaller radius than the expected radius of the learnt inhibitory connections (found empirically). This further implies that we do not start the simulation with random lateral connections, but with already well-formed Gaussian profiles, with spread $\sigma_e^{(0)}$ and $\sigma_i^{(0)}$. This has the other advantage of making the self-organization process easier and faster.

These simplifications are not in contradiction with what we want to achieve. Indeed, our purpose is not to show that self-organization can yield a one-dimensional mapping from the first to the second layer (which has already been demonstrated – in even more complex situations – in [44]), but to study the implications of a peaky prior for the tuning properties of the neurons in the second layer.

4.3.3 Simulation

The parameters for the simulation are the same as in the SOM model. They were given previously, page 30.

The parameters that are specific to the LISSOM-like model are given in the table below.

Parameter	**Value**
$\sigma_e^{(0)}$	20
$\sigma_i^{(0)}$	25
γ_a	0.5
γ_e	0.5
γ_i	0.499
η_a	0.001
η_e	0
η_i	0.0001

It may be useful to recall that the simulation is done in two steps. First, the system organizes and settles to a stable state, using a flat distribution of input stimulus. This state is called "state before learning", even though it is obtained after "training". Second, the system is kept in its final state, and further trained with a peaky prior of input, which represents the long exposure to one single orientation, characteristic of psychophysical experiments. The final state is called

"state after learning".

4.3.4 Results

The weight patterns before and after learning are reported in figure 4.6. The afferent weights show the same kind of distortion as in the SOM model. Other changes follow naturally from this geometrical reorganization, as explained page 31: the weight "blobs" amplify near the TO, as they get sharper. As to the lateral interactions, they are also modified by learning. The reorganization of the afferent weights impacts on the correlation of the activities in the second layer. Let us consider one neuron. Since more neurons are now coding for the TO and around, and since – due to the input distribution – they are more often activated, this neuron is thus getting more strongly correlated with neurons closer to the TO than him, compared to before learning. As a result, the lateral inhibitory connections will also get the same distortion as the afferent weights (it can been checked, although for the sake of clarity we do not plot the figure here). The excitatory connections being fixed, the combined lateral connections end up like in figure 4.6 (bottom right).

The resulting tuning curves and variabilities were computed as they were in the SOM model (see page 29). In figure 4.7, we plot the tuning curves after learning. Of course, the tuning curves before learning are uniform, with peaks evenly distributed along the x-axis. The shifts in orientation preference are more modest than in SOM (maximum shift is 8 deg). From the right plot, we observe a clear sharpening of the tuning curves of neurons tuned at and around the TO, and a similar broadening for neurons tuned further away.

4.3.5 Comparison to physiological data

As we did for the SOM model, we compare the changes in the slopes of the tuning curves, at the trained orientation, before and after learning, with the pattern of changes found in A. Schoups' data. This time, we get a very good correspondence between our model and physiology (figure 4.8, top row). This correspondence is even better as far as the signal-to-noise ratio is concerned (bottom row).

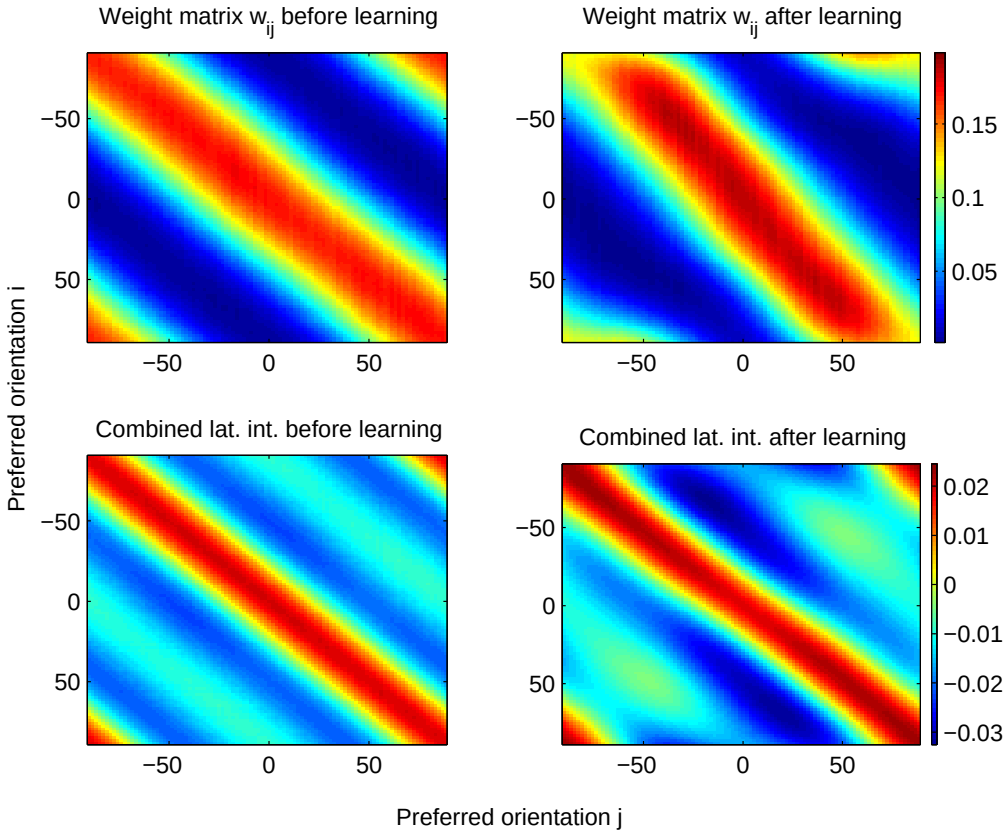

Figure 4.6: *Patterns of weights before (***left***) and after (***right***) learning with a peaky prior, in the LISSOM-like model. Two plots on the same row share the same color scale, indicated on the right.* **Top row***: afferent weights* $\mathbf{w_a}$ *connecting the first and the second layer.* **Bottom row***: the combined lateral interactions. Computed as* $[\gamma_e \mathbf{w_e} - \gamma_i \mathbf{w_i}]$.

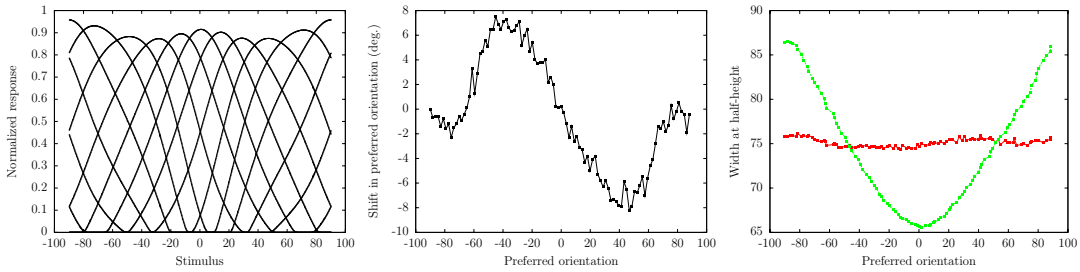

Figure 4.7: **Left***: tuning curves after learning.* **Middle***: the shift in preferred orientation for each neuron (identified by its preferred orientation on the x-axis).* **Bottom left***: the width of the tuning curve of each neuron, before learning (red) and after learning (green).*

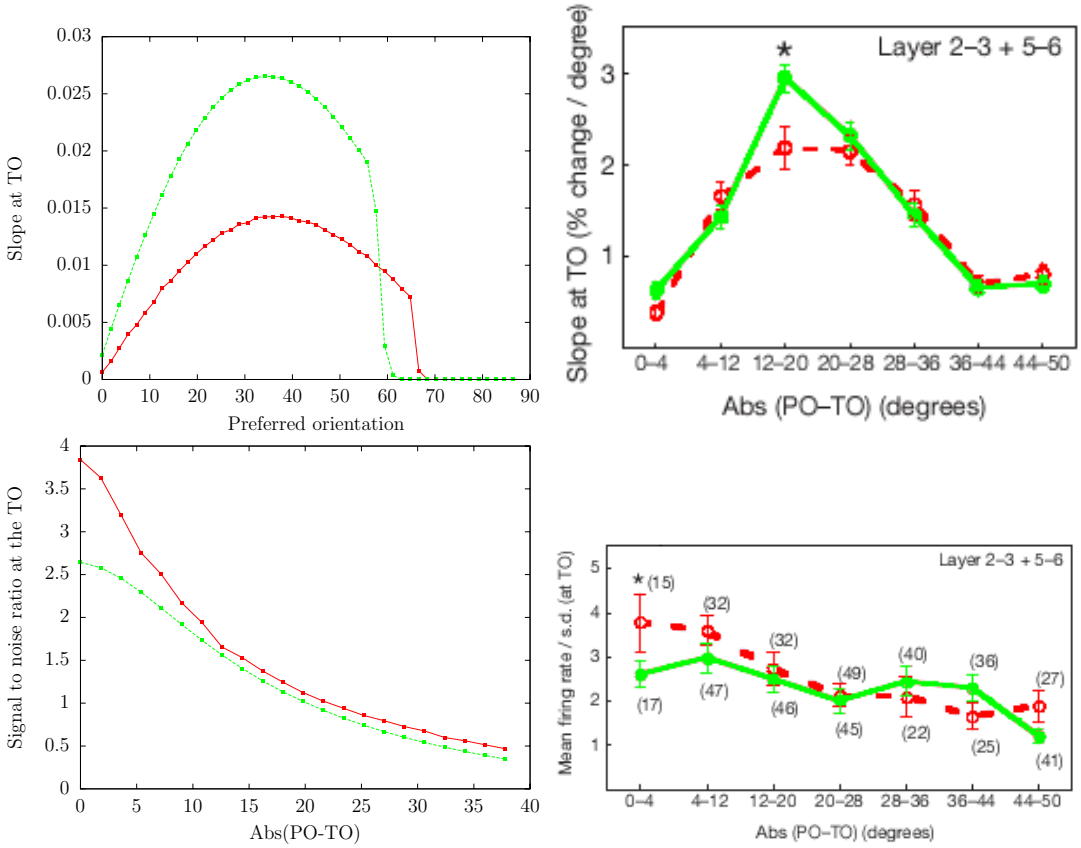

Figure 4.8: **Top left**: *the slope at the trained orientation of the tuning curve of each neuron, in our simulation, before learning (red) and after learning (green). Each neuron is represented by the distance of its preferred orientation to the TO (Abs(PO-TO)).* **Bottom left**: *the signal-to-noise ratio of each neuron* at the trained orientation. **Top and bottom right**: *the same data obtained by Schoups* et al. *from physiological recordings. Reprinted (and recolored) from [56].*

4.3.6 Comparison to another model of perceptual learning

The tuning curve we obtain after learning are interestingly very similar to the tuning curves obtained by A. Teich and N. Qian who used completely different methods [65]. They used a recurrent model of orientation selectivity, in which the sharp orientation tuning was ensured by lateral connections resembling a Mexican-hat. There were short-range excitatory connections as well as long-range inhibitory ones. They were also able to compute tuning curves from this model. Tuning curves "before learning" were computed once the network had settled to a stable state, with difference-of-Gaussian lateral connections, uniform across neurons. "Learning", in their model, consisted in slightly reducing the level of excitation to neurons tuned at and around the TO. The resulting tuning curves "after learning" are depicted in figure 4.9, where our tuning curves are plotted against them for comparison. They obtained a similar pattern of shift in orientation preference, as well as a similar pattern of sharpening and widening of the tuning curves. In addition, in our model, the tuning curves of neurons tuned around the TO "bend" toward the TO, becoming asymmetrical in addition of becoming sharper, as they do in Reich's model. Similarly, we observe that neurons tuned a bit further away have tuning curves that extend toward the TO, getting broader.

Our model does not replicate, however, the gain depression found in [65]. Rather, we find a slight gain amplification for neurons tuned at and close to the TO, as well as for neurons tuned at and close to the orthogonal orientation. In between, there is a small gain depression, for those "most informative" neurons with highest slope at the TO.

The similarity between our results and that of A. Teich can be explained quite straightforwardly, by noticing that the distortion of the lateral inhibitory connections we get have similar effects, in terms of the resulting combined lateral interactions, than a reduction of excitation around the TO. Indeed, in figure 4.6 (bottom right), we see that the excitatory part of the weights for neurons tuned around the TO is reduced.

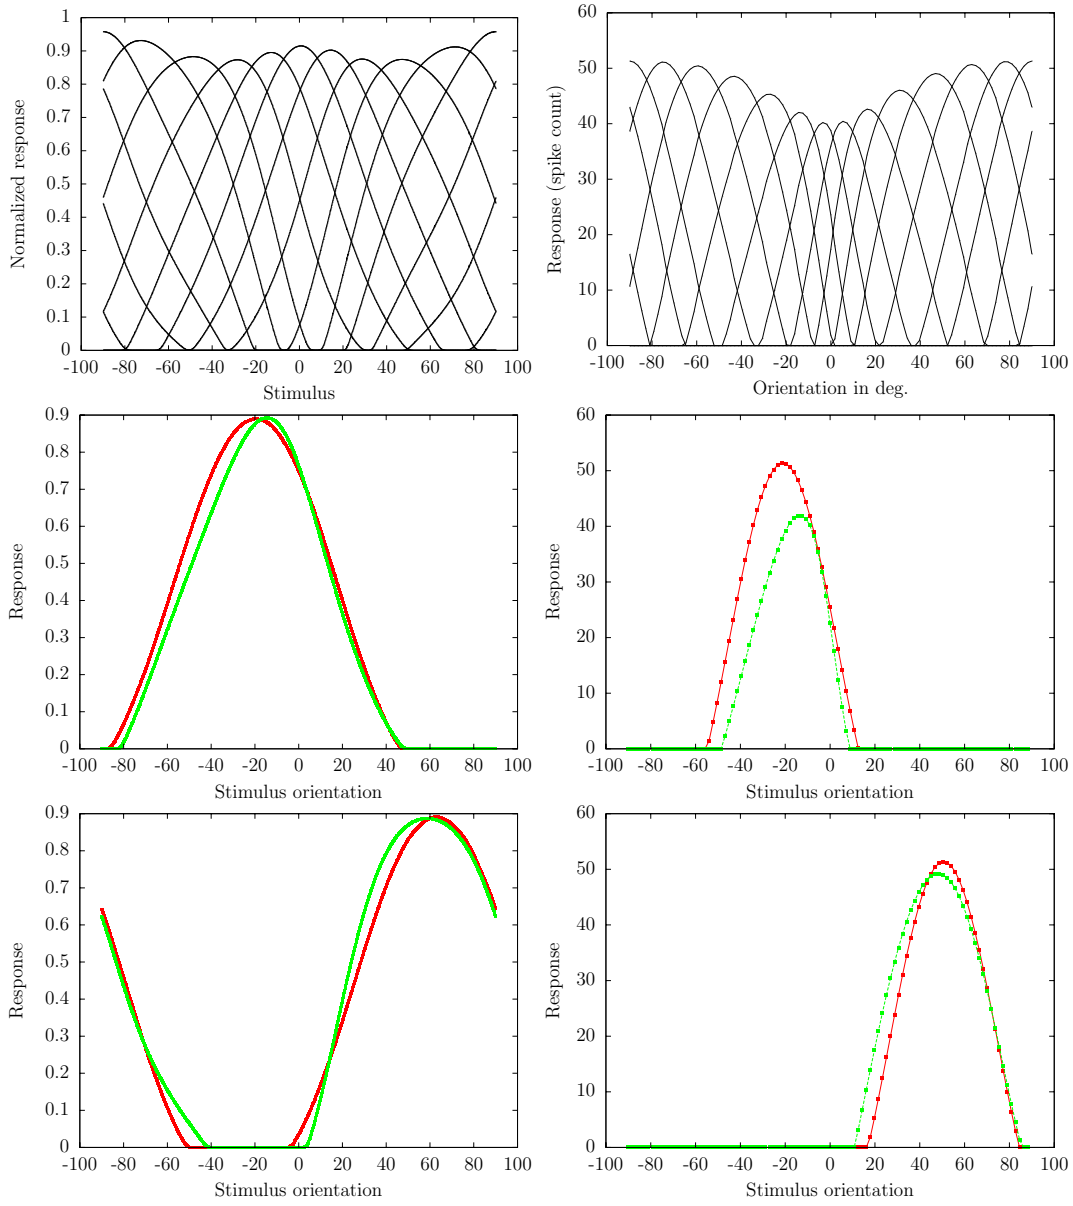

Figure 4.9: *Comparison of the tuning curves after learning in our LISSOM-like model (***left***) and in Teich's model (***right***). It should be mentioned that the tuning curves before learning were similar (though ours were a bit broader) in both models, so that this figure allows a comparison of the* patterns of changes *induced by learning. Right plots reproduced from Teich's matlab code.*

4.3.7 Discussion

Our simulation shows that it is possible to encode the stimulus prior probability through fine re-tuning of the neuronal response properties of the population. An unsupervised learning mechanism – where long-range lateral connections are capable to learn – applied to a peaky prior of input offers results that are more consistent with previous physiological findings than in a similar model with fixed lateral interactions.

Nonetheless, there are a couple of feature in our model that deserve to be further discussed. First, why did our model learn only inhibitory connections? In the literature, the alteration of long-range connections is thought to be responsible for a wide range of perceptual abilities and behavioral phenomena. As reviewed in [44], they may mediate visual comparisons (object recognition, segmentation, ...), perceptual filling-in (compensating for blind spots, perceptual completion and illusory contours), as well as the tilt-aftereffect or the brightness-contrast illusion. It seems reasonable that a phenomenon like perceptual learning, so close in nature to these kinds of illusions, would stem from similar modifications of the long-range connectivity. C. D. Gilbert advocates further in [23], when speculating on the nature of the neural changes underlying perceptual learning:

> *The dynamic changes in visual perception resulting from context and from experience do not operate independently, but rather show a close interdependence. The visuotopic extent of lateral interactions can be increased by training, as observed by the facilitation in the visibility of a target line by a colinear line. This observation, as with the receptive field plasticity seen with real and artificial scotomata, may involve a potentiation in the strength of the long-range horizontal connections.*

Second, the systematic adaptation of the activation function – so as to fit the range spanned by the feed-forward contribution to the activity in the second layer – is not well-founded. Although it happened to yield the "best" results in terms of realistic predictions of fine re-tuning, it has been argued that a more biologically plausible mechanism would be the homeostatic adaptation of neuronal excitability described by J. Triesch in [66]. We worked on this hypothesis as well, in the project, although we chose to keep the results for this discussion. For the reader not to get bogged down, we postpone the description of its implementation to appendix B. In the context of homeostatic adaptation, neurons adapt their transfer function so as to fire, in average, at a fixed rate μ. The mechanism works

properly in the first phase of our simulation: when the input distribution is even, all the neurons end up having the same activation function, and their activities are kept within non-saturating part. In the second phase (learning), when the inputs follow a peaky distribution, neurons tuned at and around the TO are naturally more often activated, therefore it become harder and harder to activate them (because they have to keep their average firing rate constant). This results in a gain depression for those neurons, and it is difficult to control the importance of this depression. If the prior is not made sharp enough, other effects become hardly noticeable, just because there is not much difference between learning and not learning. If, on the contrary, the prior is very sharp, the depression strongly dominates over other effects such as sharpening of the tuning curves or shift in orientation preference. In fact, this could be, in itself, a possible mechanism for perceptual learning, and needs to be further investigated.

If the literature of unsupervised learning has long focussed on dramatic map reorganizations, we have shown here that more subtle alterations of the tuning properties can also occur with such learning mechanisms. It does not only include "geometrical" modifications of the tuning curves, but also a specific improvement of the signal-to-noise ratio at the trained stimulus. From this study, we conclude that the repeated exposure to one particular stimulus attribute during the course of learning a perceptual task is very likely to be crucial for the establishment of long-term perceptual improvement embodied by long-term neural modifications. It sounds very challenging, however, to come up with novel psychophysical experiments that would demonstrate this importance. They should show that when a task is learnt that involves a particular orientation, and when the context is such that the subject is exposed to evenly distributed orientations, then no or little perceptual improvement is achieved. Such a setting would probably involve distractors at random orientations to make up for the trained orientation, but there would be other issues related to attention and context dependence.

Chapter 5

From neural activity to behavioral predictions

5.1 Introduction

In the previous chapters, we have set up a model of stimulus encoding, and we have studied the extent to which an unsupervised learning mechanism can modify this encoding. The following two chapters are rather independent from this first part. We want to bridge physiology and psychophysics. How much behavioral improvement can be predicted given the neural changes reported in [56, 21, 68, 49]? To carry out such a computational study, we need a model of behavioral prediction, that is, a model in which the activity of the population in response to a stimulus can be processed to assess what is perceived. One intuitive path toward the prediction of perception is that of *population decoding*. In this scheme, the stimulus is explicitly estimated from the population response. Although it seems far from obvious (and even very unlikely) that, in achieving a discrimination task, the involved brain areas explicitly need the value of the stimulus at some point, it seems rather intuitive, however, that the more precise the estimations are, the better the discrimination abilities. We therefore use this decoding approach to assess the perceptual improvement yielded by the neural modifications. In this chapter, we describe the methods used in the next chapter.

5.2 Population decoding

Given a vector of response $\mathbf{r}$ (a trial) such as that of figure 3.2 page 23, the stimulus θ responsible for this joint response is to be estimated. Several methods can do this [48]. We briefly review the population vector decoder, before giving more details about maximum likelihood (ML) and maximum a posteriori (MAP) decoders.

These decoders are so-called "estimators" from a statistical viewpoint. As such, they have two major properties: bias and variance. Imagine one single stimulus presented repeatedly, eliciting population responses and thus estimations. The bias is the difference between the average estimation and the true value of the stimulus. The variance is the standard quantity used to assess the fluctuations of the estimations around their mean, which somehow represents the confidence the estimator has in telling its estimations. A perfect estimator would be the constant function giving the true value of the stimulus each time. The bias and variance would be simultaneously null. However, for real estimators, bias and variance are subject to the so-called "bias-variance dilemma": usually, one cannot expect to have both low bias and variance. As we shall see in the following (section 5.3), the derivatives of the bias with respect to the stimulus also play an important role.

5.2.1 Population vector decoding

The certainly most obvious way to assess θ from $\mathbf{r}$ is to compute a sum of the preferred orientations of all neurons, weighted by the corresponding firing rates. Formally, this would give

$$\hat{\theta}_{\mathrm{PV}}(\mathbf{r}) = \sum_{i=1}^{N} r_i \theta_i$$

However, the stimulus belongs to a class of real variables called "circular variables". Indeed, in our case, orientation is distributed on half a circle. If we compute the standard mean, and if neurons with preferred orientations $-89°$ and $89°$ (thus differing from only $2°$) have evoked the same response, their weighted sum would be $0°$, i.e. the perpendicular orientation! Therefore, we have to do

circular statistics in order to compute this weighted mean.

$$\hat{\theta}_{\text{PV}}(\mathbf{r}) = \text{circular\_mean}\left\{(\theta_i, r_i), i = 1 \dots N\right\} \tag{5.1}$$

Practically, the circular mean is assessed by

- rescaling the θ_i's to the full circle ($\theta_i \leftarrow 2\theta_i$)
- computing $c = \sum r_i \cos(\theta_i)$ and $s = \sum r_i \sin(\theta_i)$
- if $c > 0$ and $s \geq 0$, then $\hat{\theta}_{\text{PV}} \leftarrow \arctan\left(\frac{s}{c}\right)$
- otherwise, if $c = 0$ and $s > 0$ then $\theta_{\text{PV}} \leftarrow \frac{\pi}{2}$
- otherwise, if $c < 0$ then $\hat{\theta}_{\text{PV}} \leftarrow \pi + \arctan\left(\frac{s}{c}\right)$
- otherwise, if $c \geq 0$ and $s < 0$ then $\hat{\theta}_{\text{PV}} \leftarrow 2\pi + \arctan\left(\frac{s}{c}\right)$
- otherwise $\hat{\theta}_{\text{PV}} \leftarrow 0$
- finally, $\hat{\theta}_{\text{PV}} \leftarrow \frac{1}{2}\hat{\theta}_{\text{PV}}$ to make up for the initial rescaling.

This method was originally developed by Georgopoulos and colleagues to infer the direction of arm movement in monkey, from neural population recordings [18]. Later on, it was used to decode population activity in the pre-motor cortex, the parietal area 5 and the cerebellum. In the domain of vision, it has been used to decode the responses of parietal neurons coding for the direction of motion of an object in the visual field, as well as the responses of inferotemporal neurons selective to faces (see the introduction of [52] for further references).

As we shall see in the next sections, there are some other more sophisticated methods requiring more detailed information about the response of the coding neurons. In the case of population vector, the only thing we need to know about the neurons is the locations of the peaks of their tuning curves. It has been shown that PV performs as accurately as the more complex methods when the tuning curves are best approximated by cosines [52, 64]. In practice, since tuning curves are usually well fitted by Gaussians, the wider the tuning curves, the closer they are to cosines, so the closer is the performance of PV to that of other optimal methods described below.

5.2.2 Maximum Likelihood decoding

Another possible way of estimating the stimulus that evoked response $\mathbf{r}$ is to choose

$$\hat{\theta}_{\mathrm{ML}}(\mathbf{r}) = \underset{\theta}{\operatorname{argmax}}\, p(\mathbf{r}|\theta)$$

so that the likelihood of the response given the stimulus is maximized. In practice, we maximize the log-likelihood, that is

$$\hat{\theta}_{\mathrm{ML}}(\mathbf{r}) = \underset{\theta}{\operatorname{argmax}}\, \log\left[p(\mathbf{r}|\theta)\right] \tag{5.2}$$

which leads to a computationally more stable algorithm. Indeed, no analytical expression of $\hat{\theta}_{\mathrm{ML}}$ can be derived (at least without any approximation), and the optimization procedure is done numerically (e.g. gradient ascent). Since the likelihood is, in its simpler form as we shall see below, a product of N small individual probabilities, it is a very small positive function itself, with very low gradients. Working on its logarithm thus makes things easier. Furthermore, [30] underlined that the log-likelihood can be approximated by a sum of the activity-dependant contributions of all neurons in the population, thus more likely to be implementable in the neural hardware than a product.

To perform ML decoding, we need to be able to compute the likelihood of a response $\mathbf{r}$ given a stimulus θ. The conditional likelihood of a vector depends on each individual terms $p_i(r_i|\theta)$, but also on the dependencies between each neurons. In the general case, no analytical formula is available. Within the (strong) assumption of noise independence, however, we simply have

$$p(\mathbf{r}|\theta) = \prod_{i=1}^{N} p_i(r_i|\theta)$$

where each individual $p_i(r_i|\theta)$ is given by one of equations 3.4 and 3.5.

An example of such log-likelihood as a function of the stimulus is depicted in figure 5.1. In general, provided there are enough neurons in the population, it is a well-behaved function that peaks near the real stimulus value. We see that the sharper the curvature of the likelihood is, the more confident we are in telling the estimated value of the stimulus. Indeed, if the likelihood is very curved, potential stimuli a bit further away from the peak are very quickly "dismissed" because of evoking lower and lower likelihood values for the response vector $\mathbf{r}$. The reader

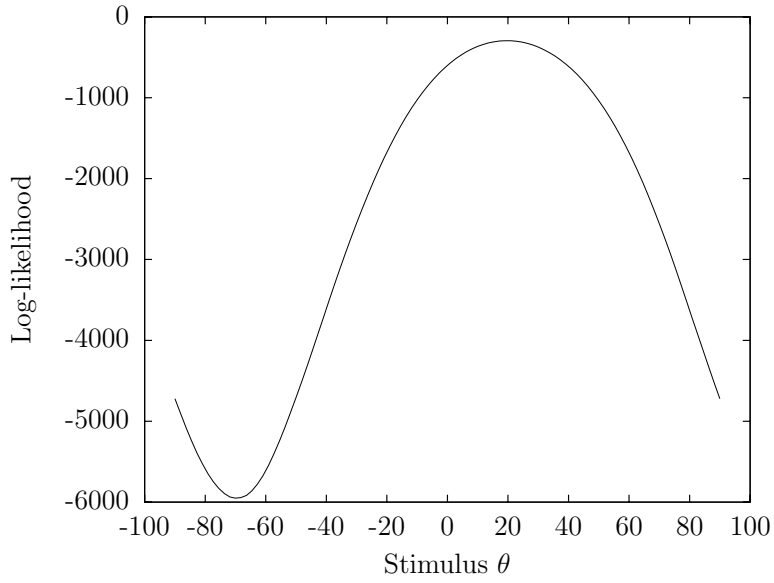

Figure 5.1: *Log-likelihood $p(\mathbf{r}|\theta)$ as a function of θ, where $\mathbf{r}$ is the trial plotted in figure 3.2 (so the real stimulus was 20 deg). This curve would slightly change with different trials for the same stimulus. The population has 100 neurons with preferred orientations evenly distributed on half a circle, the noise is Gaussian with Fano factor set to 1.3. Tuning curves were the same as in figure 3.1. For information, the ML procedure applied to this trial gives $\hat{\theta}_{ML}(\mathbf{r}) = 19.692 \pm 10^{-3}$.*

will like to relate this back to equation 3.6 defining Fisher information as the expected value of this log-likelihood curvature.

A very important property of the ML estimator is that, in the limit of large populations of neurons, it is *unbiased* and *its variance saturates the Cramer-Rao bound*. Thus, if there are enough neurons in the population, the variance of the ML estimator is well approximated by the inverse Fisher Information. In this respect the ML decoder is said to be optimal. In fact, there is another asymptotic condition for ML to be optimal: that of large spike counts in the time window used to measure the rates.

5.2.3 Maximum A Posteriori decoding

The previous estimator, ML, maximizes the likelihood without taking the prior probability of the stimulus into account. If we think of perceptual learning as a process modifying the encoding of a "prior expectation" on the stimulus, then we need to use a decoding scheme that incorporates this alternative.

One way to use the information from the prior is to maximize the posterior probability defined as the product of the likelihood and prior. This is known as Maximum A Posteriori (MAP). We have :

$$\hat{\theta}_{\text{MAP}}(\mathbf{r}) = \underset{\theta}{\text{argmax}}\, p(\mathbf{r}|\theta) \cdot p(\theta) \tag{5.3}$$

In practice, again, we maximize the logarithm of the above quantity, so that

$$\hat{\theta}_{\text{ML}}(\mathbf{r}) = \underset{\theta}{\text{argmax}}\ (\log\left[p(\mathbf{r}|\theta)\right] + \log\left[p(\theta)\right]) \tag{5.4}$$

The difficulty lies in estimating the prior $p(\theta)$, since it is far from obvious what it looks like before learning, and how it may be affected by numerous presentations of the same orientation. The main idea, in our studies, is to build a prior that is severely biased toward a rather narrow area around the trained orientation.

5.3 Simulation of psychophysical experiments

5.3.1 Introduction

We have seen that behavioral improvement with practice has been observed for different psychophysical tasks. Here, I review the different methods used to derive the relevant measures of performance that I have already introduced throughout the text: Just-Noticeable-Difference (JND), discriminability (d'), percent correct, decision bias.

The underlying theory is called "Signal Detection Theory" (SDT). To compute these measures of performance, we must bear in mind that the subject has to *make a decision* from a perceived signal that is a stochastically encoded version of the real stimulus, thus carrying much *uncertainty*. Figure 5.2 sketches out the path leading from the stimulus to the decision. Here, we are interested in the latest stage (decision making). In order to simulate the decision and assess performance, we have to make an assumption on the strategy in play. SDT provides different strategies that are optimal with respect to the different tasks, from a probabilistic viewpoint.

Here, we are mostly interested in the two types of tasks previously mentioned in the text: one- and two-intervals tasks. In both, the subject must choose between two answers. They differ, however, in the number of times (called "intervals") in which the subject is presented a stimulus before having to decide. I will directly apply the theory to orientation discrimination, and show how we can relate the properties of our estimators (namely their bias and variance) to behavioral performance.

Nota This is given more as a reference than as a preliminary background for the rest of the dissertation. A quick look at equations 5.6, 5.8, 5.10 and 5.11 would be enough to carry on the reading of chapter 6. Most of the material presented in this section and in appendix B was found in [37].

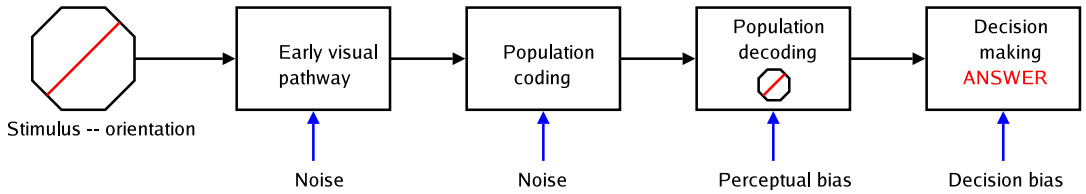

Figure 5.2: *From the stimulus to the decision. The last stage is the only one that depends on the type of psychophysical task. In this model, there are two sources of bias. The perceptual bias is the statistical bias of the estimator used at the population decoding stage. The decision bias is related to the strategy used in the decision making process. In the psychophysics literature, one may find the adjectives "liberal" or "conservative", depending on whether the subject tends to favor high hit rates (see in the text) or low false-alarms rates.*

5.3.2 Preliminary definitions

z-transform Imagine that we have a Gaussian distribution $f(s)$ with mean μ and standard-deviation σ. Any value s on the s-axis can be given a "z-score"

$$z_s = \frac{s-\mu}{\sigma}$$

representing the distance of s from the mean μ, in standard-deviation units. Now assume we are given such a z-score z. We want to know the probability that a stimulus s drawn from the same distribution has a z-score z_s less than z. This is merely the area under the normal distribution function $\mathcal{N}(0,1)$ from $-\infty$ to z. This is only dependent on z, and defines a p-transform, function of z. Simple integration manipulations give the following formula:

$$p(z) = \frac{1}{2}\mathrm{erfc}\left(-\frac{z}{\sqrt{2}}\right)$$

This function is monotonically increasing from 0 to 1, therefore it has an inverse, which is called the $z-transform$, extensively used in SDT.

$$z(p) = -\sqrt{2}\ \mathrm{erfcinv}(2p)$$

Hits and false-alarms As mentioned above in the introduction, both task families we are interested in here are about choosing between two alternatives. For

the sake of simplicity and without loss of generality, let us call them "Yes" and "No". In the vocabulary of SDT, a *hit* corresponds to reporting "Yes" while the correct answer is "Yes", and a *false alarm* amounts to reporting "Yes" whereas the correct answer was "No". From these definitions are defined the *hit rate* H and *false alarm rate* F, corresponding to the frequencies at which the subject generates a hit of a false alarm. Note that, in general, $H + F \neq 1$.

Discriminability A major measure of sensitivity of the observer, in SDT, is the discriminability d'. It is defined as

$$d' = z(H) - z(F) \tag{5.5}$$

where H and F are defined above, and z is the z-transform. When the observer cannot discriminate at all, and when he/she is unbiased (does not favor hits nor false alarms), $H = F$ and the discriminability is null. When the observer is optimal, d' tends to infinity. In practice, this never happens, and d' drops rapidly as H increases and F decreases a bit. For example, $d' \simeq 4.65$ when $F = 0.01$ and $H = 0.99$.

A case of interest for us is the following. Discrimination involves making a difference between the estimates that come from orientation θ_1 and those coming from θ_2. In other words, the discriminability is a measure of how separated the two distributions of estimates are. Let us assume that $\theta_1 < \theta_2$. It can be shown that, in the majority of cases, the distributions of estimations are Gaussians[1]. To set things, let us say that stimulus θ_1 evokes estimates following a Gaussian distribution with variance σ_1^2 and mean $\theta_1 + b_1$ where b_1 is the bias of the estimator. The same for the estimates of θ_2, with σ_2^2 and b_2. Is can be shown that d' is very well approximated by

$$d' = \frac{(\theta_2 + b_2) - (\theta_1 + b_1)}{\sqrt{\frac{\sigma_1^2 + \sigma_2^2}{2}}}$$

When $\theta_2 - \theta_1$ is small enough, $\sigma_1^2 \simeq \sigma_2^2 \simeq \sigma^2$. Furthermore, d' can be reformulated

[1]This is in fact true when the estimates come from ML decoding (figure 5.3), and in the limit of large populations. But in this dissertation, we assume that whatever the decoder, the estimate ends up being normally distributed (fair approximation in practice).

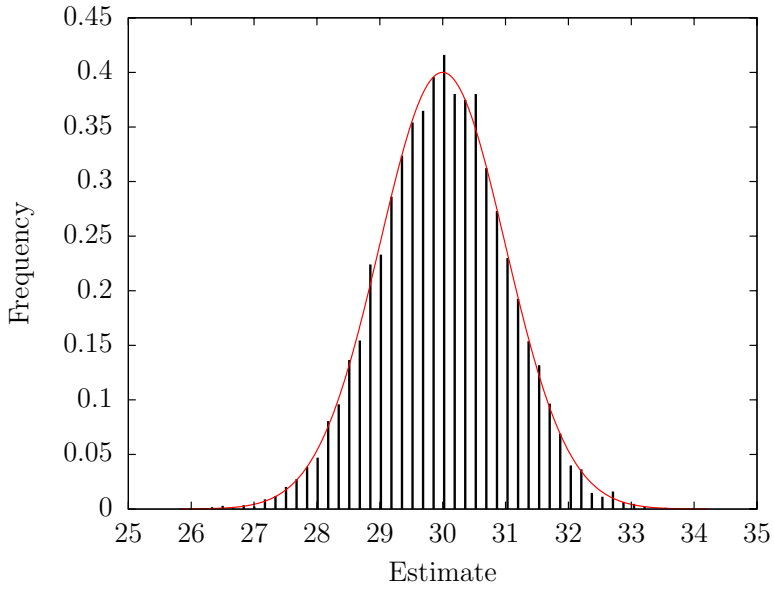

Figure 5.3: *Histogram for ML estimates. We simulated 10,000 trials. The stimulus was 30 degrees. The population had same parameters as in previous "toy experiments" such as that of figure 3.2. In red, the Gaussian distribution with same mean and variance is plotted for comparison. We see that the ML estimates approximately follow a Gaussian distribution, centered on the real value of the stimulus (ML is unbiased).*

as

$$\begin{aligned} d' &= \frac{(\theta_2 - \theta_1)}{\sigma}\left(1 + \frac{(b_2 - b_1)}{\theta_2 - \theta_1}\right) \\ &= \frac{\Delta\theta \cdot (1 + b')}{\sigma} \end{aligned} \tag{5.6}$$

Decision bias Imagine a radiologist trying to detect tumors from MRI reconstructed images. The consequences of reporting "there is a tumor" while there is actually no tumor are not disastrous. On the contrary, reporting "there is no tumor" while there is one is a serious mistake. Therefore, if in doubt, the doctor will prefer to report the presence of a tumor, thus deliberately raising both the false-alarm rate F and the hit rate H. The subject's willingness to say "yes" is measured by the decision bias. In SDT, it is defined as:

$$c = -\frac{1}{2}\left[z(H) + z(F)\right] \tag{5.7}$$

In the rest of this chapter, we always assume that the subject has no decision bias.

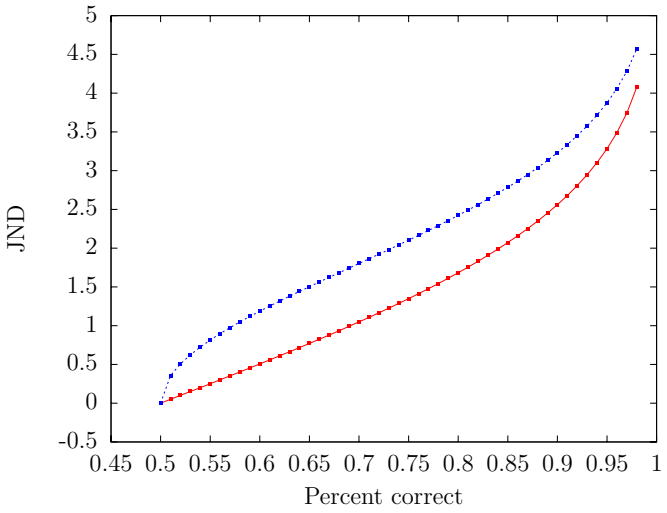

Figure 5.4: *JND as a function of its corresponding percent correct $p(c)$. Red: one-interval task; Blue: two-intervals task. The estimator is assumed to have unit variance and null bias, to fix things.*

Percent correct It is merely the frequency $p(c)$ at which the observer gives a correct answer. It depends on d' in a way that is specific to the type of task as we shall see later. It obviously depends on the difficulty of the task, that is, the difference between the two stimuli to be discriminated.

Just Noticeable Difference A commonly accepted measure of the performance in orientation discrimination is the *Just Noticeable Difference*, defined as the orientation difference (between the two stimuli to be discriminated) that elicits a given percent correct (usually 84%). As mentioned previously, $p(c)$ depends on the discriminability. Therefore, in order to compute the "JND 84%", we have to find $d'_{84\%}$ corresponding to $p(c) = 0.84$, and then to invert equation 5.6 to obtain:

$$\mathrm{JND}_{84\%} = \frac{\sigma \cdot d'_{84\%}}{1 + b'} \tag{5.8}$$

The dependency between the JND and the percent correct is shown in figure 5.4 for both task paradigms. The underlying equations follow in section 5.3.3.

Here, we see that **behavioral performance is dependant on the bias and the variance of the estimator we use, only**. More specifically, we will see that a high variance is often counterbalanced by a repulsive bias ($b' > 0$). In the literature, researchers have sometimes tried to assess performance with biased

estimators, without taking b' into account, which obviously lead to inaccurate results.

From equation 3.10 (Cramer-Rao bound) we see that the JND is bounded by:

$$\mathrm{JND}_{84\%} \geq \frac{d'_{84\%}}{\sqrt{I_F}} \tag{5.9}$$

where I_F is the Fisher information.

5.3.3 How to relate $p(c)$ and d'?

The link between the percent correct and the discriminability depends on the task. For each task, SDT provides an optimal strategy within which one can mathematically derive the formula linking the two quantities. Derivations of the following two equations are given in appendix A.

One-interval task

$$p(c) = \frac{1}{2}\mathrm{erfc}\left(-\frac{d'}{2\sqrt{2}}\right) \tag{5.10}$$

Two-intervals task

$$p(c) = \left[\frac{1}{2}\mathrm{erfc}\left(-\frac{d'}{2\sqrt{2}}\right)\right]^2 + \left[\frac{1}{2}\mathrm{erfc}\left(\frac{d'}{2\sqrt{2}}\right)\right]^2 \tag{5.11}$$

Chapter 6

Simulation of neural changes and comparison to psychophysics

6.1 Introduction

In our attempt to bridge neurophysiological and psychophysical experiments, we want to answer the question

> *To what extent can the observed neural changes account for the behavioral improvement?*

As we saw at the end of chapter 2, neural correlates of perceptual learning in the early stages of visual processing are not yet well understood. Contradicting data makes it difficult to answer the above question in a purely quantitative fashion. Besides, there may not be only one way to model the tuning curves after learning. Data is often focusing on what happens to the tuning curves *around* the trained orientation (e.g. how does the slope change? the signal-to-noise ratio? ... etc), giving very few clues about how tuning is affected elsewhere by learning. Therefore, we are going to work on *idealized* models of neural changes, and investigate what are their implications in terms of behavioral predictions. We apply decoding methods to assess performance and compare our results to psychophysics.

First, we implement a simple sharpening model where tuning curves are merely made narrower around the trained orientation. Second, we implement a simple gain-modulation model where tuning curves are either positively or negatively

modulated in amplitude around the trained orientation. Finally, we revisit a model of modified tuning curves that [65] came up with by changing relevant parameters in a recurrent model of orientation selectivity. This latter modelling combines sharpening, gain modulation, shift in orientation preference, and asymmetries of the tuning curves. It has the chief advantage that it fits the physiological data found in [56] while having strong biological foundations. Furthermore, we saw in chapter 4 that the same patterns of changes can be obtained in a different model based on unsupervised Hebbian learning, which gives it even more credit.

For each of these models, we compare performance between trained and naive populations. As we mentioned before in the introduction, a major theoretical question is that of where learning actually occurs in the discrimination process: is it only the encoder that changes, or does the decoder change as well ? The available physiological data makes it straightforward to research the effects of a modification of the encoding part. Indeed, within our model, the stimulus is primarily encoded in the neural responses (tuning curves and variability), which is precisely what is found to change in physiology. To relate these changes to behavior, however, we can either use a decoder that does not change with learning, or a decoder that adapts as the subject trains. The first type of decoder is embodied by the PV procedure, which – provided there are no shifts in the preferred orientations of the neurons – remains unchanged with learning. The second category is represented by ML decoding, where the likelihood of a set of responses is modified by learning. As we will see below, PV and ML predict approximately the same amount of behavioral improvement, and it turns out that predictions in terms of learning fall short of expectation when compared to psychophysics.

6.2 Simple sharpening model

A simple way to model the change in slope at the T0 found in [56] and [49] is to sharpen the tuning curves of neurons tuned around the T0, so that they get steeper at the T0. We keep any other parameter constant.

Practically, the tuning curve of neuron i is given a width at half-height W_i fol-

lowing "narrowing profile" :

$$W_i = W_i^{(0)} \cdot \left[1 - A \ \exp\left(-\frac{(\theta_i - \theta_{\text{TO}})^2}{2\sigma_w^2}\right)\right] \tag{6.1}$$

where A determines the percentage of narrowing, and σ_w its spread. $W_i^{(0)}$ is the width at half height of the tuning curve before learning, and θ_i is the preferred orientation of neuron i.

In our simulations, tuning curves were Gaussian, noise was Gaussian with a constant Fano factor. Here is the set of parameters (left) and the widths of the tuning curves after learning according to equation 6.1 (right) :

Parameter	**Value**
N	100
Baseline	10 spikes
Amplitude	50 spikes
Fano factor	1.3
$W_i^{(0)}$	70 deg. for all i
A	0.4
σ_w	20 deg.
θ_{TO}	$\approx$ 20 deg.

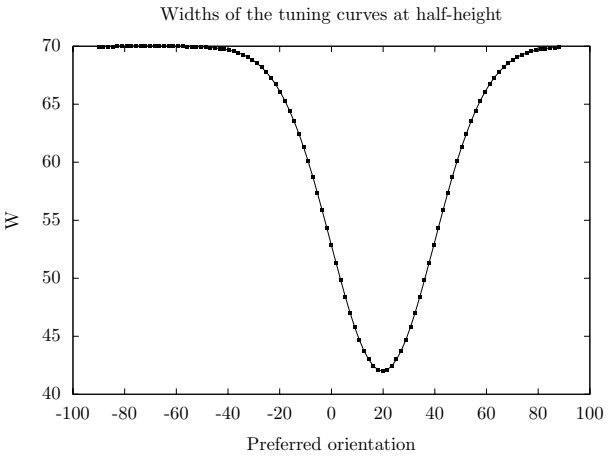

Using these parameters, the slope at TO changes in roughly the same way as observed in [56], as reported in figure 6.1. By "roughly" I mean that neurons termed "most informative neurons" (those with highest Fisher information at TO) are not actually the same before and after learning, but tend to move towards neurons tuned closer to the TO, which can be seen in figure 6.1 (right). This pitfall is addressed by the model of [65] where more complex changes in the tuning curves are modelled.

We start investigating how performance is affected by learning under the assumption that the stimulus is decoded using Population Vector (PV) or Maximum Likelihood (ML) decoders. As we saw, performance is only dependant on the bias and the variance of the estimator. Figure 6.2 shows these properties for both estimators, as a function of the test orientation. Both are unbiased before learning. For ML, this is natural since it is an unbiased estimator for large populations.

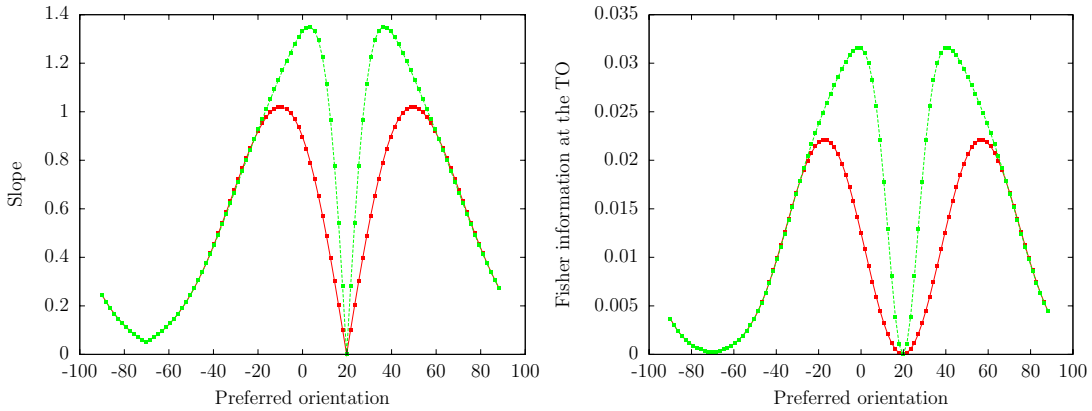

Figure 6.1: **Left:** *slopes at the trained orientation for each neuron. Neurons are referred to using their preferred orientation.* **Right:** *Fisher information at TO for each neuron. Red: before learning, Green: after learning.*

For PV, this is due to the fact that the population is initially homogeneous, circularly symmetric. After learning, however, PV becomes strongly biased. Its bias is repulsive around the TO since its slope is positive. Therefore, although the variance gets worse around the TO, performance is compensated by the bias. To understand this, one can think of equation 5.8: a repulsive bias increases the denominator $(1 + b')$, which in turn decreases the JND. Similarly, one can think in terms of separability of estimation distributions: given two stimuli $s_1 < s_2$, the means of the two distributions of estimations are all the farther from each other as the bias is steeply increasing (the bias at s_1 being then much less than that at s_2). The variance being locally flat (compared to the bias), the resulting separability $\frac{\mu_2 - \mu_1}{\sigma}$ increases, lowering the JND.

Performance before and after learning is reported in figure 6.3. Here, ML is not significantly better than PV because tuning curves are broad ($W_i^{(0)} = 70$ deg.) which means that they can be well approximated by cosines, thus making PV nearly optimal. Performance increases by about 24% at the TO for both estimators. In psychophysical experiments, performance has been seen to increase by more than 80% in monkey, and around 70% in human.

One could argue that, in our simulations, the initial level of performance is already high (JND = 2 deg., which is quite close to the asymptotic performance reached by well trained human beings), which could mean that there be "no room for improvement" anymore. However, we ran the same simulations with only 30 neurons (as opposed to 100): the initial JND was consistent with the initial JND

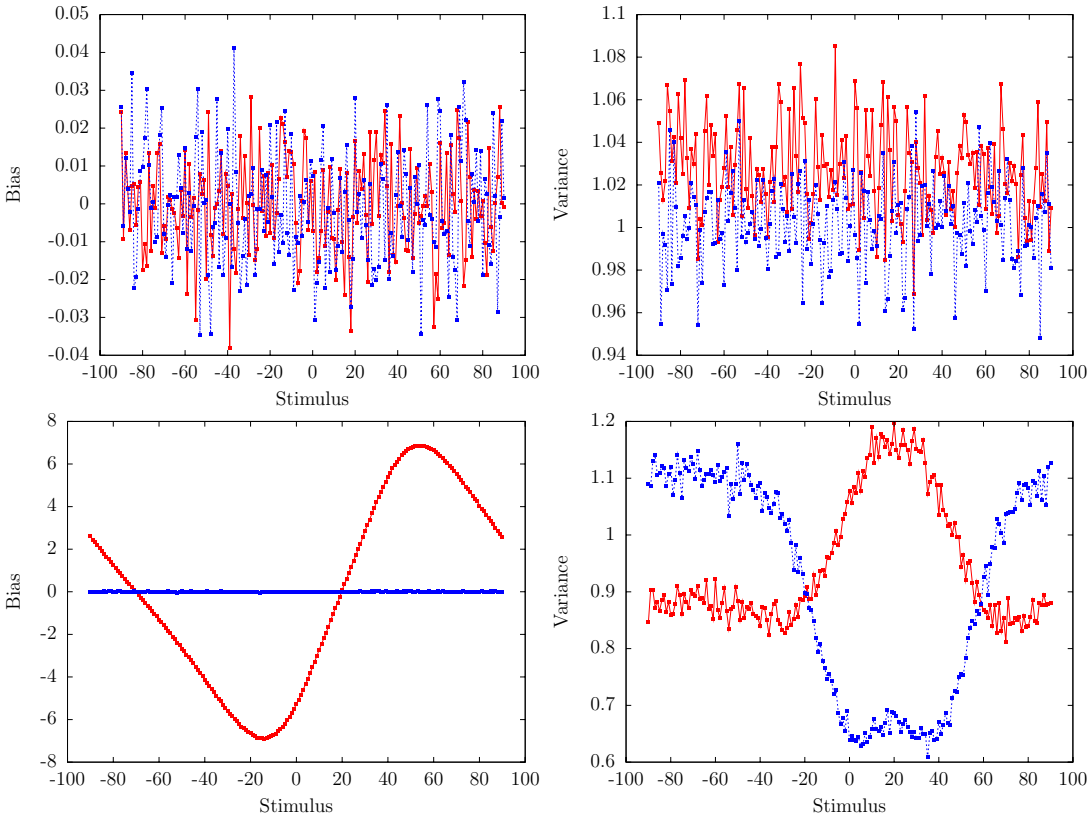

Figure 6.2: *Bias (left column) and variance (right column) of PV (red) and ML (blue) estimators, before (top line) and after (bottom line) learning, as a function of the stimulus being decoded. Note that the scale of the top plots is small, so that the biases can in fact be considered as constant, not noisy.*

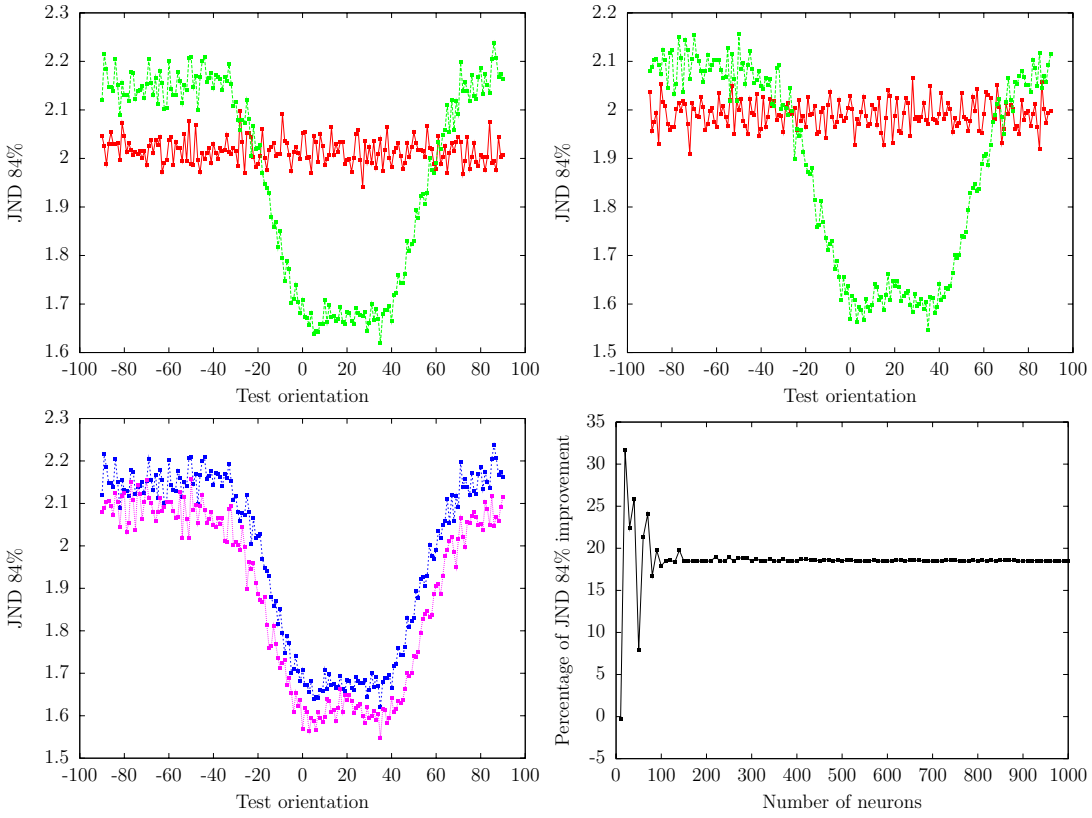

Figure 6.3: *Psychophysical performance (JND) under the assumption of PV (top left) and ML (top right) decoding. Red: before learning. Green: after learning. The bottom-left plot compares performance of both estimators after learning (blue: PV, purple: ML). The bottom-right plot reports the amount of learning (as described by the percentage of JND improvement, computed using "asymptotic AI ML") as a function of the number of neurons in the population. It is clearly invariant when the population is large enough.*

in human (about 5 degrees), and the amount of improvement was of the same order of magnitude than in previous experiments (22%). To be more precise, the percentage of JND improvement is plotted as a function of the number of neurons in the population (see figure 6.3, bottom right), and we notice that is is asymptotically (and quickly) independent of N. The percentage of improvement is defined as

$$p = \frac{\mathrm{JND}_{\mathrm{before}} - \mathrm{JND}_{\mathrm{after}}}{\mathrm{JND}_{\mathrm{before}}}$$

From equation 5.8 we see that, for an unbiased estimator, this fraction reduces to:

$$p = \frac{\sigma_{\mathrm{before}} - \sigma_{\mathrm{after}}}{\sigma_{\mathrm{before}}}$$

Using the Cramer-Rao bound, and assuming that the estimator is optimal, we conclude that

$$p \simeq \frac{\left(\frac{1}{\sqrt{I_F}}\right)_{\mathrm{before}} - \left(\frac{1}{\sqrt{I_F}}\right)_{\mathrm{before}}}{\left(\frac{1}{\sqrt{I_F}}\right)_{\mathrm{before}}}$$

From equations 3.7 or 3.8 it can be shown that Fisher information varies linearly in N. Therefore, the $\frac{1}{\sqrt{N}}$ terms in the fraction above cancel, hence the invariance. This invariance is reassuring in the sense that real brain populations of neurons involved in a discrimination task are likely to be much larger than the small populations we model herein.

However, this invariance does not apply to the absolute performance (the JND itself instead of its improvement): as mentioned above, a population of only 30 neurons can replicate the baseline performance observed in human before learning. Conversely, a model with a biologically realistic population size would yield a totally unrealistic JND (something like 1/1000 degree). This suggests that some other factors would make the task harder in reality, cancelling the power of neural pooling. A very probable one is the presence of correlations between neurons, not modelled in this dissertation, and that could possibly decrease the number of "degrees of freedom" among neurons, reducing the model's predicted JND to that of 30 independent cells.

From figure 6.3 it is also worth noticing that, at the orientation orthogonal to the TO, performance is worse after learning than before. This result has been observed in [56] (although it fell short of statistical significance for one of their three human subjects).

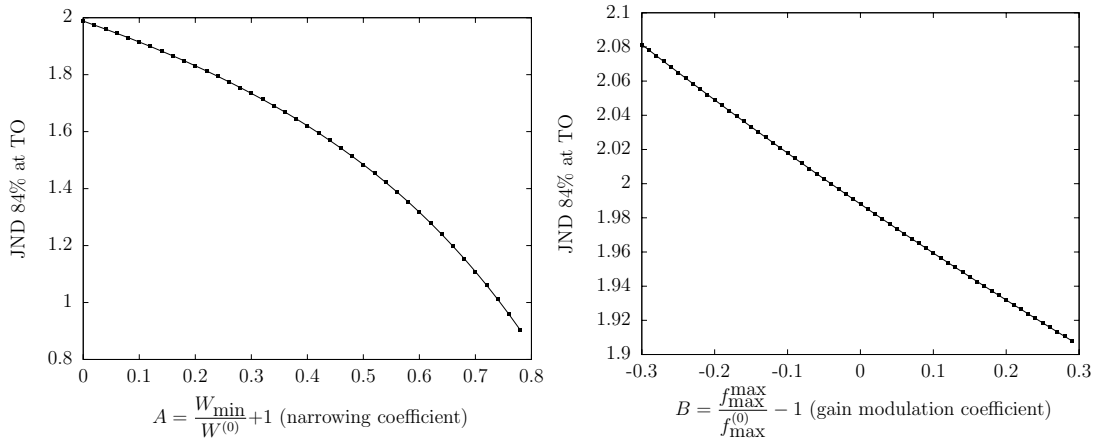

Figure 6.4: *Performance at the TO as a function of the narrowing coefficient A (left) or the gain modulation coefficient B (right) used to build the population after learning. For each A or B, a new population is built in which the widths of the tuning curves follow equation 6.1 or 6.2, and the JND is computed using the "optimal ML procedure" (assuming that Cramer-Rao bound is saturated). Note that the scales are different, so the right plot is actually flatter than the left one.*

Our experiments show that it is difficult to explain a 80% JND improvement by only sharpening the tuning curves and using PV or ML decoding, and this, independently on the number of neurons in the population. To make it even more conclusive, we report in figure 6.4 the JND with increasing narrowing coefficients (A in equation 6.1). We see that in order to reach a 50% JND improvement, tuning curves have to be made more than 70% narrower around the TO, which is not realistic according to neurophysiological data ($A = 0.7$ means that slopes at TO for most informative neurons are multiplied by 4, whereas [56] found they were multiplied by no more than 1.5).

6.3 Simple gain modulation

Physiological data sometimes report a specific gain modulation of the tuning curves around the trained orientation. These data are controversial: [55] and [20] report that the response amplitude of neurons tuned around the TO *decreases* in V1, whereas [49] found an *increase* in V4. In this section, we research the effects of positive and negative gain modulation. As we did in the previous section,

performance is assessed using PV and ML.

By gain modulation, we mean that the maximum firing rate of the neuron is raised, but the baseline remains the same, and the tuning curve is still Gaussian. Thus, the tuning curve of neuron i is given a maximum firing rate $f_i^{\max}$ following "gain modulation profile" :

$$f_i^{\max} = f_i^{\max,(0)} \cdot \left[1 + B \ \exp\left(-\frac{(\theta_i - \theta_{\text{TO}})^2}{2\sigma_w^2}\right)\right] \tag{6.2}$$

where B determines the percentage of amplification or depression (following its sign), and σ_w its spread. $f_i^{\max,(0)}$ is the width at half height of the tuning curve before learning, and θ_i is the preferred orientation of neuron i. As in the narrowing model, we took $\sigma_w = 20$ deg. in all our experiments.

We first apply the PV decoder to assess the perceptual improvement when learning performs a 20% gain amplification ($B = 0.2$), and when it does a 20% gain depression ($B = -0.2$). The results are reported in figure 6.5.

First, both amplification and depression – when achieved within the bounds of physiological realism (not more than 30%) – yield very little change in performance. Second, amplification and depression have opposite effects on the bias and variance of the estimator, as well as on the resulting performance: amplification helps, depression harms.

To explain why gain modulation has a much lower impact on performance than narrowing the tuning curves, we can relate it back to equation 3.7, that gives Fisher information as proportional to:

$$I_F(\theta_{\text{TO}}) \propto \sum_i \frac{f_i'^2(\theta_{\text{TO}})}{f_i(\theta_{\text{TO}})}$$

Narrowing the tuning curves of neurons tuned around θ_{TO} decreases the denominator in any case, and so makes the code more accurate. It also increases the numerator in some cases (if the narrowing profile is well-chosen), which makes the effect even more remarkable. If, however, learning uses gain modulation to maximize the TO, the result is less obvious. Decreasing $f_i(\theta_{\text{TO}})$ by gain depression causes the numerator $f_i'^2(\theta_{\text{TO}})$ to decrease even more (power of two), and the resulting performance is then worse than before learning. Better result may be achieved by using gain amplification, since the numerator then increases more than the denominator, thus slightly enhancing the code accuracy.

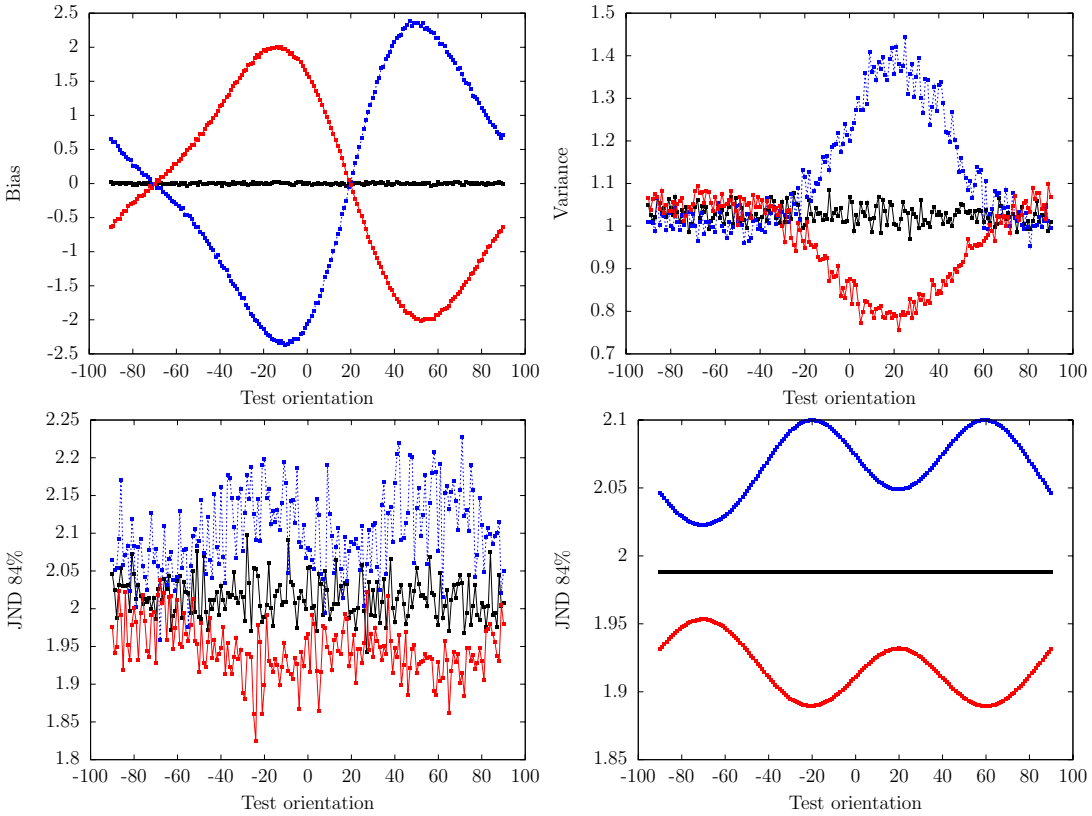

Figure 6.5: *Black: before learning; Red: gain amplification ($B = 0.2$); Blue: gain depression ($B = -0.2$) – Bias (top left) and variance (top right) of PV estimator. Behavioral performance is then assessed from them (bottom left). For comparison, performance computed with Fisher information (ML in the limit of large N) is plotted (bottom right).*

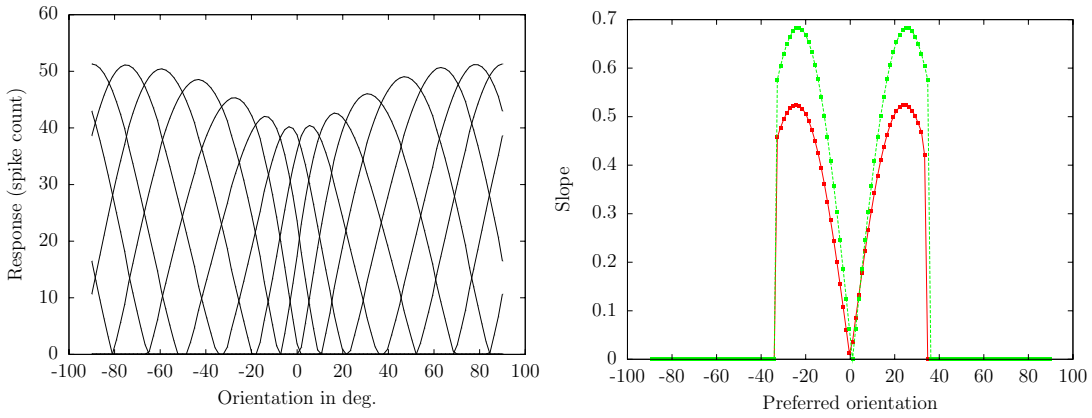

Figure 6.6: **Left:** *a few tuning curves in the population after learning, taken from [65].* **Right:** *the slope of the tuning curves at the trained orientation for all neurons (each neuron being identified with its preferred orientation). Red: before learning; Green: after learning.*

As we did for the narrowing model, we report the predicted discrimination performance as a function of B, the gain modulation coefficient (see figure 6.4). We see that a reasonable gain modulation (from 10% to 30%, according to [20] and [55]) cannot account for the actual discrimination improvement in human or primate. We shall now combine narrowing and gain amplification, using the tuning curves in [65].

6.4 Revisiting Teich & Qian, 2003 [65]

In their paper, A. F. Teich and N. Qian studied perceptual learning as a mechanism that modifies the connection properties in a recurrent model of V1 orientation selectivity. More specifically, they showed that slightly reducing the net excitatory connections to cells tuned around the TO leads to the same patterns of tuning curves as observed in physiology (figure 6.6, left). In particular, an increase in the slope at the TO is found for those most informative neurons only, consistent with the data in [56] (figure 6.6, right).

This model incorporates many different patterns of changes of the tuning curves: narrowing and gain depression around the TO, shift in the preferred orientations toward the TO, but also an asymmetry of the tuning curves for neurons tuned around the TO, not specifically reported by physiological recordings. The tuning

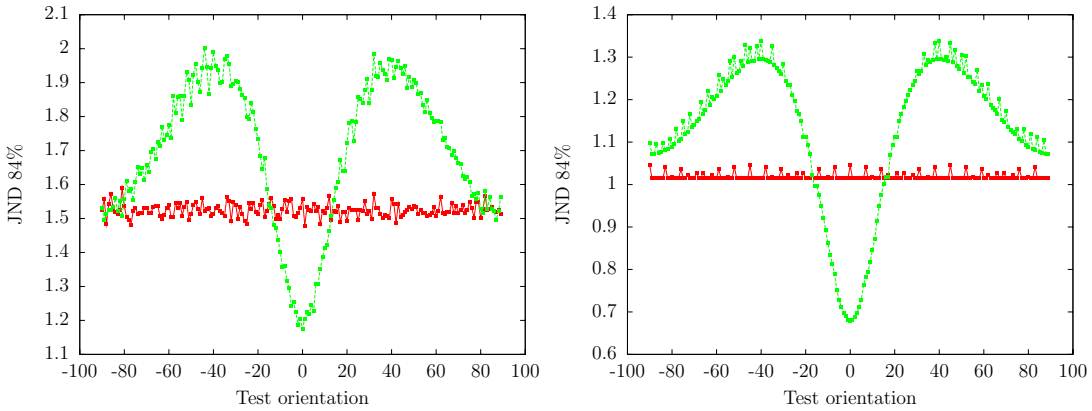

Figure 6.7: *Performance in orientation discrimination before and after learning, when stimulus is decoded using PV (left) and ML (right).*

curves actually "bend toward the TO" (which implies the peak shift and the asymmetry) while trying to keep their positions on the other side. It looks like a compromise between increasing the slope of the tuning curves (hence, the Fisher information) at the TO, and not over-covering the area around the TO.

We report the predicted behavioral performance in figure 6.7 for both PV and ML decoders. The maximum JND improvement is obtained with the optimal ML decoder, and reaches about 33%, which is significantly higher than what we found in previous models. With PV, only 22% is predicted. Thus, combining multiple neuronal effects – that have been shown to yield some perceptual improvement (previous sections) separately – produce an even larger behavioral effect. Still, in any case, it stays much below the behavioral improvement reported in psychophysics.

It should be mentioned that the assumption of noise independence that we make in order to compute the behavioral predictions for sure does not hold in Teich's model, where neurons are laterally interconnected, thus being necessarily correlated.

6.5 Discussion

We have used very simple models of neural modifications in order to relate physiology and psychophysics. We have shown that as far as the representation of the stimulus can change so as to remain physiologically realistic, it is impossi-

ble to explain the dramatic perceptual improvement characteristic of perceptual learning in vision. Our results are consistent with previous studies (usually done by physiologists in the same papers that describes their data) based on the idea of "ideal Bayesian observer", in which they also compared psychometric curves obtained from their data.

An alternative to learning the encoder would be to learn the decoder. In the course of this project, we have investigated the effects of a prior probability of stimulus on the predictions of the MAP decoder (see page 54). We noticed a "scaling issue" which, as far as we know, has never been reported in the literature. Let us have a look at equation 5.3, page 54. On the one hand, the term $p(\theta)$ is typically less than 1, since $\int p(\theta)$ should be unit. On the other hand, the likelihood $p(\mathbf{r}|\theta)$ is of many orders of magnitude smaller, as a product of small quantities (equation 3.1). Therefore, the likelihood strongly dominates, and our simulations – not reported in details here – showed no difference between ML and MAP in this context. For the prior to have an actual effect on the maximum of the posterior, is has to be "rescaled" to match the order of magnitude of the prior. Even when rescaled, the prior does not really make a difference with ML in terms of predicted behavioral improvements. In fact, the rescaled prior attracts the stimulus that maximizes the posterior toward the trained orientation, but an attractive bias is bad for discrimination purposes. The negative effects of the prior was compensated by a lower variance around the TO. Hence, we found that this alternative was not really relevant.

Another direction would be to contrast the performance of a globally suboptimal decoder and of a locally optimal one, and assess the resulting behavioral improvement. Pioneering work in this direction has been carried out by H. S. Seung, G. Mato, and H. Sompolinsky [58, 38]. They described several linear decoders (the "fully adaptative perceptron", the "vector discriminator") capable of adapting to a particular stimulus (e.g. trained orientation) and compared them with the suboptimal but uniformly performing ML and PV estimators. They reported "a tradeoff between performance at the adapted stimulus and the degree of transfer" to other stimuli.

Similarly, a whole theory of perceptual learning as a process of selective reweighting as been elaborated recently by A. Petrov, B. A. Dosher and Z.-L. Lu [46]. In their framework, learning occurs only in the "read-out" connections

to a decision unit, while the stimulus representation remains unchanged. Contrary to the hypothesis we made in chapter 4, namely that of perceptual learning as a mechanism modifying the lateral interactions, Petrov's theory relies on an "incremental associative updating of the projection between different areas, as opposed to updating the lateral connections within a specific area".

Appendix A

Derivation of equations 5.10 and 5.11

A.1 One-interval task

In a "one-interval task", only one stimulus is presented in each trial. There is no reference offered for comparison[1]. In the studies we have seen, this kind of task corresponds to the "clockwise - counterclockwise" experiments. In each trial, the stimulus rotated clockwise or counterclockwise from the standard orientation is presented (at random), but the reference itself is never presented again in the trial. After seeing the stimulus, the subject has to tell either CW or CCW.

Imagine that we are to simulate the decision process. The information we have is the estimate of the stimulus orientation, $\hat{\theta}_s$. As we said in 5.3.2, when the stimulus presented is CW (resp. CCW), this estimate follows a Gaussian distribution centered in μ_{CW} (resp μ_{CCW}), with variance σ^2_{CW} (resp σ^2_{CCW}). Formally, making a decision from this estimate amounts to determining which of these two conditional distributions $\hat{\theta}_s$ is most likely to have been drawn from. This kind of decision can be made on the basis of various criteria. [26] considers four decision goals : maximizing the proportion correct, maximizing a weighted proportion of hits and correct rejections, maximizing the expected value according to some reward/punishment criterion, and testing a statistical hypothesis. Each of these 4

[1] Actually, this claim is very controversial within the context of training, where feedback is given after each trial, thus providing implicit information about the reference.

goals can actually be achieved using the same underlying test called the *likelihood ratio* test. In this test, given an estimate $\hat{\theta}_s$ of the stimulus, the likelihood ratio

$$\text{LRatio} = \frac{p(\hat{\theta}_s|\text{ccw})}{p(\hat{\theta}_s|\text{cw})}$$

is compared to a *criterion* k. If LRatio $< k$ then the decision is CW, otherwise the decision is CCW. In the case of two Gaussians, this test is obviously equivalent to a simple decision boundary θ_t: if $\hat{\theta}_s < \theta_t$ then we decide CW, otherwise CCW.

It can be shown that if the subject has no decision bias, then $k = 1$, and

$$\theta_t = \frac{\mu_{\text{cw}} + \mu_{\text{ccw}}}{2}$$

In which case, the *percent correct* measure of performance is maximized (first of the previously mentioned goals). For what follows, we will make this assumption of unbiased decision. The percentage correct is simply derived as follows:

$$p(c) = p(\text{cw}) \cdot p(\hat{\theta}_s < \theta_t|\text{cw}) + p(\text{ccw}) \cdot p(\hat{\theta}_s > \theta_t|\text{ccw})$$

where $p(\text{cw})$ and $p(\text{ccw})$ are the prior probabilities of the two stimuli. Using the basic properties of a Gaussian distribution, we find that

$$p(c) = \frac{1}{2}\text{erfc}\left(-\frac{d'}{2\sqrt{2}}\right) \qquad \text{with} \qquad d' = \frac{\mu_{\text{ccw}} - \mu_{\text{cw}}}{\sigma}$$

Note that this is not dependent on the prior probabilities anymore. This would not be the case if the decision maker were biased, thus pushing the criterion θ_t leftward or rightward. By inverting equation 5.10, we are able to assess d' correponding to, say, $p(c) = 0.84$, which gives the $\text{JND}_{84\%}$. For an estimator with unit variance and null bias, the JND is related to its corresponding percent correct as depicted by the red curve in figure 5.4.

A.2 Two-intervals task

In a two-intervals task, a first stimulus is presented in the first time interval, followed by another stimulus in the second time interval. The task might then be for the observer to tell the order in which he saw the stimuli, or to tell whether the stimuli were identical or not, for example.

Here I will focus on the second task, namely "match-to-sample" or "same-different". It is used, for example, in [68]. There are only two different stimuli that can occur: either an orientation tilted clockwise from 45 deg. (tilt Δs), either the counterclockwise equivalent. Therefore, there are four possible sequences:

$$(\text{ccw}, \text{ccw}) \qquad (\text{cw}, \text{cw}) \qquad (\text{cw}, \text{ccw}) \qquad (\text{ccw}, \text{cw})$$

Now imagine we are to simulate the decision of the observer. The observer can make his choice from the estimates of the first stimulus $\hat{\theta}_1$ and of the second stimulus $\hat{\theta}_2$. These estimates are drawn from two conditional distributions $p(\theta|\text{cw})$ and $p(\theta|\text{ccw})$, approximately Gaussians with means μ_{cw} and μ_{ccw} and same variance σ. Again, we assume that the observer is unbiased with respect to his choice. Therefore, he will tell that the stimuli were identical if both estimates $\hat{\theta}_1$ and $\hat{\theta}_2$ lie on the same side of the boundary $\frac{1}{2}(\mu_{\text{cw}} + \mu_{\text{ccw}})$. On the contrary, he will tell they were different if they lie on opposite sides of the boundary.

A simple derivation of the percent correct gives:

$$p(c) = \left[\frac{1}{2}\text{erfc}\left(-\frac{d'}{2\sqrt{2}}\right)\right]^2 + \left[\frac{1}{2}\text{erfc}\left(\frac{d'}{2\sqrt{2}}\right)\right]^2$$

The fact that it does not depend on the prior probabilities of "same" and "different" stimuli comes from the fact that the subject is unbiased. They would matter if the subject had a tendency to favor hits or to avoid false-alarms.

Again, equation 5.11 can be numerically inverted in order to compute the $\text{JND}_{84\%}$. In order to do that, we just need to use the fact that $\text{erfc}(-x) = 2 - \text{erfc}(x)$, so that equation 5.11 becomes

$$p(c) = 1 - \text{erfc}(x) + \frac{1}{2}\text{erfc}^2(x)$$

which is a simple quadratic equation in $\text{erfc}(x)$, with $x = -\dfrac{d'}{2\sqrt{2}}$. The only valid solution is

$$\text{erfc}\left(-\frac{d'}{2\sqrt{2}}\right) = 1 + \sqrt{2p(c) - 1}$$

which can be indeed inverted using the inverse erfc function.

Appendix B

Homeostatic learning rule

As an alternative to the piecewise linear function used in our LISSOM-like model (chapter 4), a smoother sigmoid function can be used to keep the activity in the second layer within the (0,1) range. Moreover, instead of having one single activation function for all neurons, it is more realistic to give each neuron its own transfer function σ_i. Finally, the activation functions learn as well from the activity of the neurons.

We choose to implement a form of homeostasis of the neurons' firing rate levels given in [66]. For neuron i, the activation function is parameterized:

$$\sigma_i(x) = \frac{1}{1 + \exp\left[-(a_i x + b_i)\right]}$$

The idea is to update a_i and b_i so as to achieve an approximately exponential distribution for each neuron's firing rate. Cortical recordings have indeed revealed that cortical cells exhibit such an exponential distribution of firing rate in response to natural images. Besides, it can be shown that this distribution has the highest entropy among all distributions of a non-negative random variable with a fixed mean [2], which optimizes the ratio between the energy spent in firing and the amount of information being transmitted.

J. Triesch [66] showed that if x is the pre-synaptic activity and $y = a_i x + b_i$ is the resulting post-synaptic firing rate, then a_i and b_i must be updated according

to the following equations:

$$\begin{aligned} a_i &\leftarrow a_i + \varepsilon_h \left[\frac{1}{a_i} + x - \left(2 + \frac{1}{\mu}\right) xy + \frac{1}{\mu} xy^2 \right] \\ b_i &\leftarrow b_i + \varepsilon_h \left[1 - \left(2 + \frac{1}{\mu}\right) y + \frac{1}{\mu} y^2 \right] \end{aligned} \tag{B.1}$$

where ε_h is the learning rate of the homeostatic plasticity (much smaller than ε_a, ε_e and ε_i), and μ is the mean firing rate to be achieved.

Bibliography

[1] M. Ahissar and S. Hochstein. The reverse hierarchy theory of visual perceptual learning. *Trends in Cognitive Sciences*, 8:457–464, 2004.

[2] R. Baddeley, L. F. Abbott, M. Booth, F. Sengpiel, and T. Freeman. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. In *Proceedings of the Royal Society London*, number 264 in B, pages 1775–1783, 1998.

[3] K. Ball and R. Sekuler. A specific and enduring improvement in visual motion discrimination. *Science*, 218:697–698, 1982.

[4] K. Ball and R. Sekuler. Direction-specific improvement in motion discrimination. *Vision Research*, 27:953–965, 1987.

[5] G. G. Blasdel. Orientations selectivity, preference, and continuity in monkey striate cortex. *The Journal of Neuroscience*, 12:3139–3161, 1992.

[6] D. Boussaoud, R. Desimone, and L. G. Ungerleider. Visual topography of area teo in the macaque. *The Journal of Comparative Neurology*, 306:554–575, 1991.

[7] C. W. G. Clifford. Perceptual adaptation: motion parallels orientation. *TRENDS in Cognitive Sciences*, 6(3):136–142, 2002.

[8] D. R. Cox and D. V. Hinckley. *Theoretical statistics*. London: Chapman & Hall, 1974.

[9] R. E. Crist, M. K. Kapadia, G. Westheimer, and C. D. Gilbert. Perceptual learning of spatial localization: specificity for orientation, position, and context. *The Journal of Neurophysiology*, 78:2889–2894, 1997.

[10] R. E. Crist, W. Li, and C. D. Gilbert. Learning to see: experience and attention in primary visual cortex. *Nature Neuroscience*, 4(5):519–525, 2001.

[11] P. Dayan and L. F. Abbott. *Theoretical neuroscience; Computational and Mathematical Modeling of Neural Systems*. The MIT Press; Cambridge, Massachusetts; London, England, 2001.

[12] M. Fahle. Specificity of learning curvature, orientation, and vernier discriminations. *Vision Research*, 37:1885–1895, 1997.

[13] M. Fahle. Visual learning in humans. *The Journal of Vision*, 4:879–890, 2004.

[14] M. Fahle. Perceptual learning: specificity versus generalization. *Current Opinion in Neurobiology*, 15:154–160, 2005.

[15] M. Fahle and S. Edelman. Long-term learning in vernier acuity: effects of stimulus orientation, range and of feedback. *Vision Research*, 33:397–412, 1993.

[16] M. Fahle and T. Poggio, editors. *Perceptual learning*. The MIT Press, 2002.

[17] C. S. Furmanski, D. Schluppeck, and S. A. Engel. Learning strengthens the response of primary visual cortex to simple patterns. *Current Biology*, 14:573–578, 2004.

[18] A. P. Georgopoulos, A. Schwartz, and R. E. Kettner. Neuronal population coding of movement direction. *Science*, 233:1416–1419, 1986.

[19] G. M. Ghose. Learning in mammalian sensory cortex. *Current Opinion in Neurobiology*, 14:513–518, 2004.

[20] G. M. Ghose and J. H. R. Maunsell. Perceptual learning can selectively alter neural responses in primate v1. *Society for Neuroscience Abstracts*, 23:1544, 1997.

[21] G. M. Ghose, T. Yang, and J. H. R. Maunsell. Physiological correlates of perceptual learning in monkey v1 and v2. *The Journal of Neurophysiology*, 87:1867–1888, 2002.

[22] E. J. Gibson. Perceptual learning. *Annual Reviews in Psychology*, 14:29–56, 1963.

[23] C. D. Gilbert. Adult cortical dynamics. *Physiological Reviews*, 78:467–485, 1998.

[24] C. D. Gilbert, M. Sigman, and R. E. Crist. The neural basis of perceptual learning. *Neuron*, 31:681–697, 2001.

[25] J. Gold, P. J. Bennett, and A. B. Sekuler. Signal but not noise changes with perceptual learning. *Nature*, 402:176–178, 1999.

[26] D. M. Green and J. A. Swets. *Signal Detection Theory and Psychophysics*. New York: John Wiley, Inc., 1996.

[27] M. H. Herzog and M. Fahle. The role of feedback in learning a vernier discrimination task. *Vision Research*, 37(15):2133–2141, 1997.

[28] D. Hubel and T. Wiesel. Receptive fields, binocular interaction, and functional architecture in the cat's visual cortex. *The Journal of Physiology (London)*, 160:106–154, 1962.

[29] D. Hubel and T. Wiesel. Binocular interaction in striate cortex of kittens reared with artificial squint. *The Journal of Neurophysiology*, 28:1041–1059, 1965.

[30] M. Jazayeri and J. A. Movshon. Optimal representation of sensory information by neural populations. *Nature Neuroscience*, 9(5):690–696, 2006.

[31] U. Karmarkar and Y. Dan. Experience-dependant plasticity in adult visual cortex. *Neuron*, 52:577–585, 2006.

[32] A. Karni and G. Bertini. Learning perceptual skills: behavioral probes into adult cortical plasticity. *Current Opinion in Neurobiology*, 7:530–535, 1997.

[33] A. Karni and D. Sagi. Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. *PNAS*, 88:4966–4970, 1991.

[34] A. Karni and D. Sagi. The time course of learning a visual skill. *Nature*, 365:250–252, 1993.

[35] T. Kohonen. The self-organizing map. *Proceedings of the IEEE*, 78:1464–1480, 1990.

[36] Z. Liu. Perceptual learning in motion discrimination that generalizes across motion directions. *PNAS*, 96(24):14085–14087, 1999.

[37] N. A. Macmillan and C. D. Creelman, editors. *Detection theory: a user's guide*. Routledge, 2004.

[38] G. Mato and H. Sompolinski. Neural network models of perceptual learning of angle discrimination. *Neural Computation*, 8:270–299, 1996.

[39] N. Matthews, Z. Liu, B. J. Geesaman, and N. Qyan. Perceptual learning on orientation and direction discrimination. *Vision Research*, 39:3692–3701, 1999.

[40] N. Matthews, Z. Liu, and N. Qian. The effect of orientation learning on contrast sensitivity. *Vision Research*, 41:463–471, 2001.

[41] N. Matthews and L. Welch. Velocity-dependent improvements in single-dot direction discrimination. *Perception and Psychophysics*, 59:60–72, 1997.

[42] R. Miikkulainen. Self-organizing process based on lateral inhibition and synaptic resource redistribution. In T. Kohonen, K. Mäkisara, 0. Simula, and J. Kangas, editors, *Proceedings of the 1991 International Conference on Artificial Neural Networks*, pages 415–420, Amsterdam: North-Holland, 1991.

[43] R. Miikkulainen, J. A. Bednar, Y. Choe, and J. Sirosh. Self-organization, plasticity, and low-level visual phenomena in a laterally connected map model of the primary visual cortex. In R. L. Goldstone, P. G. Schyns, and D. L. Medin, editors, *Perceptual Learning*, volume 36 of *Psychology of Learning and Motivation*, pages 257–308. Academic Press, San Diego, CA, 1997.

[44] R. Miikkulainen, J. A. Bednar, Y. Choe, and J. Sirosh. *Computational maps in the visual cortex*. New York: Springer, 2005.

[45] A. J. Parker and W. T. Newsome. Sense and the single neuron: probing the physiology of perception. *Annual Reviews in Neuroscience*, 21:227–277, 1998.

[46] A. A. Petrov, B. A. Dosher, and Z.-L. Lu. The dynamics of perceptual learning: an incremental reweighting model. *Psychological Review*, 112(4):715–743, 2005.

[47] B. Pleger, H. R. Dinse, P. Ragert, P. Schwenkreis, J. P. Malin, and M. Tegenthoff. Shifts in cortical representations predict human discrimination improvement. *PNAS*, 98(21):12255–12260, 2001.

[48] A. Pouget, P. Dayan, and R. Zemel. Inference and computation with population codes. *Annual Reviews in Neuroscience*, 26:381–410, 2003.

[49] S. Raiguel, R. Vogels, S. G. Mysore, and G. A. Orban. Learning to see the difference specifically alters the most informative v4 neurons. *The Journal of Neuroscience*, 26(24):6589–6602, 2006.

[50] G. H. Recanzone, W. M. Jenkins, G. T. Hradek, and M. M. Merzenich. Progressive improvement in discriminative abilities in adult owl monkeys performing a tactive frequency discrimination task. *The Journal of Neurophysiology*, 67:1015–1030, 1992.

[51] G. H. Recanzone, C. E. Schreiner, and M. M. Merzenich. Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. *The Journal of Neuroscience*, 13:87–103, 1993.

[52] E. Salinas and L. F. Abbott. Vector reconstruction from firing rates. *The Journal of Computational Neuroscience*, 1:89–107, 1994.

[53] K. Sathian and A. Zangaladze. Tactile learning is task-specific but transfers between fingers. *Perceptual Psychophysics*, 59:119–128, 1997.

[54] A. Schoups, R. Vogels, and G. A. Orban. Human perceptual learning in identifying the oblique orientation: retinotopy, orientation specificity, and monocularity. *The Journal of Physiology*, 483(3):797–810, 1995.

[55] A. Schoups, R. Vogels, and G. A. Orban. Effects of perceptual learning in orientation discrimination on orientation coding in v1. *Invest Ophtalmol Vis Sci Suppl*, 39:684, 1998.

[56] A. Schoups, R. Vogels, N. Quian, and G. A. Orban. Practising orientation identification improves orientation coding in v1 neurons. *Nature*, 412:549–553, 2001.

[57] L. Schwabe and K. Obermayer. Adaptivity of tuning functions in a generic recurrent network model of a cortical hypercolumn. *The Journal of Neuroscience*, 25(13):3323–3332, 2005.

[58] H. S. Seung and H. Sompolinski. Simple models for reading neuronal population codes. *PNAS*, 90:10749–10753, 1993.

[59] C. Shawn and D. Bavelier. Action video game modifies visual selective attention. *Nature*, 423:534–537, 2003.

[60] Shiu and Pashler. Improvement in line orientation discrimination is retinally local but dependant on cognitive set. *Perception and Psychophysics*, 52:582–588, 1992.

[61] M. Sigman and C. D. Gilbert. Learning to find a shape. *Nature Neuroscience*, 3:264–269, 2000.

[62] R. Sireteanu and R. Rettenbach. Perceptual learning in visual search: fast, enduring, but non-specific. *Vision Research*, 35:2037–2043, 1995.

[63] J. Sirosh. A self-organizing neural network model of the primary visual cortex. Ph.D. thesis AI95-237, Department of Computer Sciences, The University of Texas at Austin, Austin, TX, 1995.

[64] H. P. Snippe. Parameter extraction from population codes: a critical assessment. *Neural Computation*, 8(3):511–529, 1996.

[65] A. F. Teich and N. Qian. Learning and adaptation in a recurrent model of v1 orientation selectivity. *The Journal of Neurophysiology*, 89:2086–2100, 2003.

[66] J. Triesch. A gradient rule for the plasticity of a neuron's intrinsic excitability. In W. Duch and al., editors, *ICANN 2005*, volume 3696 of *LNCS*, pages 65–70. Springer-Verlag, Berlin Heidelberg, 2005.

[67] R. Vogels and G. A. Orban. The effect of practice on the oblique effect in line orientation judgments. *Vision Research*, 25(11):1679–1687, 1985.

[68] T. Yang and J. H. R. Maunsell. The effect of perceptual learning on neuronal responses in monkey visual area v4. *The Journal of Neuroscience*, 24(7):1617–1626, 2004.

[69] E. Zohary, S. Celebrini, K. H. Britten, and W. T. Newsome. Neuronal plasticity that underlies improvement in perceptual performance. *Science*, 263:1289–1292, 1994.