Neural modelling of the McCollough Effect in color vision

Thesis · August 2014

DOI: 10.13140/RG.2.1.4786.0886

2 authors, including:

Giacomo Spigler
Scuola Superiore Sant'Anna
13 PUBLICATIONS **100** CITATIONS

SEE PROFILE

Neural modelling of the McCollough Effect in color vision

Giacomo Spigler (s1360784)

Master of Science

Cognitive Science (Neural Computation & Neuroinformatics)

School of Informatics

University of Edinburgh

2014

Abstract

The present work investigates the neural mechanisms underlying the McCollough Effect through the simulation of three different models of color visual systems: a dichromatic and a trichromatic one that are inspired to the anatomy of the primates' visual system and an idealised trichromatic model designed to aid the self-organization of a simulated Primary Visual Cortex to a physiologically-plausible level.

The McCollough Effect allows for a direct investigation of cortical plasticity and homeostatic adaptation, and can provide a reference for works in the cortical modelling of color perception. Within this dissertation we compare our results to previous psychological data on humans and to the relevant work in modelling the effect, effectively reproducing its the main properties.

Acknowledgements

I would like to thank all my friends, with whom I shared so many wonderful experiences, with a special thanks to Daniele and Simone, for the amazing time spent together, and to Valerio, one of my oldest friends, for the biggest laughs in the last decade.

A big thanks goes to The University of Edinburgh, for the great environment it provided, and to my amazing supervisor James Bednar for the great advices that he could always provide. A thanks also goes to Chris Ball, for the time he spent discussing and explaining his model and code.

The biggest thanks however is for my parents, who have always supported me and gave me countless opportunities to travel, study and learn, and they taught me the value of work and education.

Edinburgh, August 13^{th}

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.

(*Giacomo Spigler (s1360784)*)

Table of Contents

Chapter 1

INTRODUCTION

The McCollough Effect is a special color aftereffect [41] that makes white, oriented gratings appear to be colored [31, 16, 42, 40, 8]. The effect is shown in Fig. 1.1 : before induction, two orthogonal white gratings appear to be colorless; induction is then performed by showing colored gratings in alternation (usually red and green, which are known to produce the strongest effect) for a total time of 5 to 15 minutes. After the adaptation to the patterns, the white gratings appear to possess a slight shade of color, the complementary to the one of the induction grating with the same orientation. Notably, the effect has been observed to last for days and even months, which greatly constrains the possible underlying mechanisms [27, 38]. Another interesting aspect of the effect is its contingency to both orientation and color of the patterns. Differently oriented gratings are still perceived as colored, but the effect fades approximately linearly with the difference in the orientation of the patterns (with respect to the induction gratings), and vanishes at around 45^o.

The McCollough Effect was first studied in depth by Celeste McCollough, who used special devices to allow subjects to *cancel* the perceived color by compensating with its opponent. The amount of "correction" was used to measure the (perceived) *colorimetric purity* of the stimulus, which in turn was found to be a good measure for the strength of the ME (McCollough Effect) [31].

The special importance of the McCollough Effect is due to two main reasons. First, it requires the interaction of color and orientation processing mechanisms at some point within the visual hierarchy, which means that we can learn how different features interact in visual cortices and how they are *combined* [15]. Second, even though many simple color aftereffects are known (e.g., adaptation to colored patches), they don't last a long time. Understanding the McCollough Effect will thus probably uncover mechanisms that affect

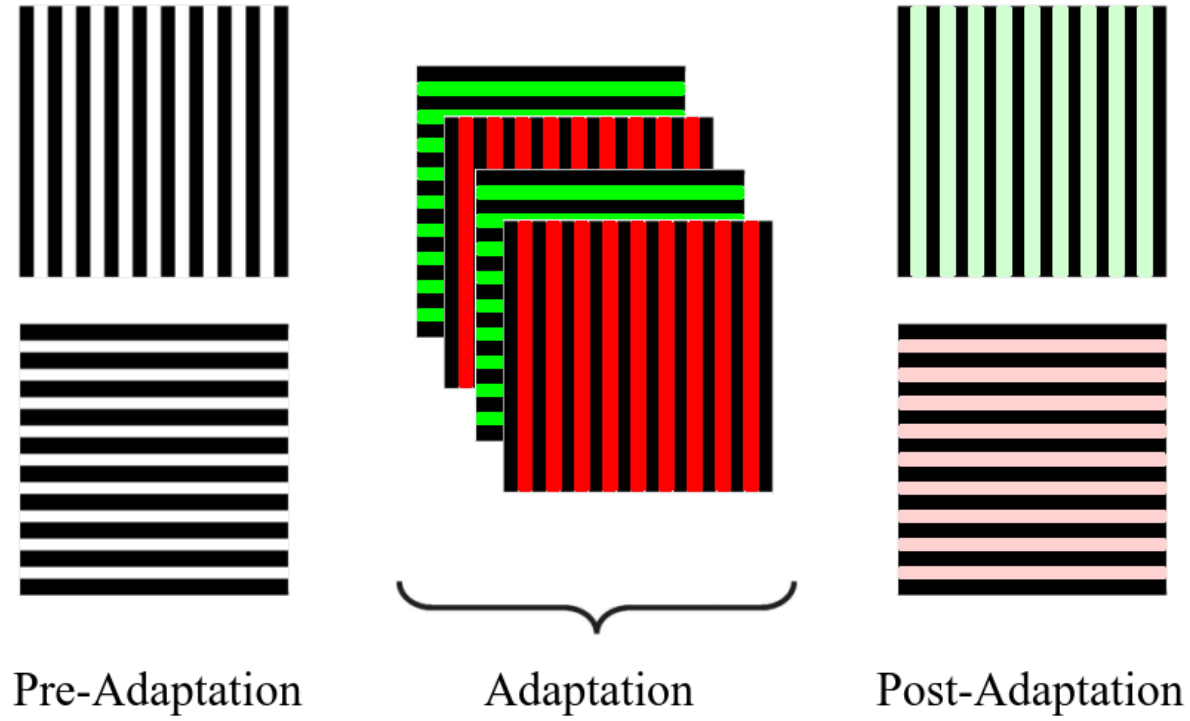

Figure 1.1: *Example McCollough Effect. colorless grids (left) are perceived as slightly colored (right) after* $5-15$ *minutes of adaptation with similar, colored patterns. The perceived color is complementary to the one of the induction grid with the same orientation.*

the long-term properties and behavior of the visual system.

However, many problems still persist. There is still no strong agreement on where is the neural locus which gives rise to the effect, and it is still not clear what are its exact mechanisms. Within the scope of our work, we will investigate the different roles of homeostatic adaptation and cortical plasticity in order to understand the dynamic adaptation of the Primary Visual Cortex to varying input statistics [45, 46]. To test our work, we will examine several adaptations of previously proposed models, which are characterized by different degrees of anatomical and functional biological plausibility.

While we are grounding our research on previous, similar, work [11], we can now make use of advanced tools that make the understanding of the system easier and more tractable.

1.1 Organization of the work

Within this dissertation we are first going to discuss the context of the problem with a broad review of the relevant background (Chapter 2). Then, we will present our models and the experiments we performed, with special focus in particular on three different models of color vision, namely a dichromatic and a trichromatic vision models, which are easy to handle and are anatomically inspired to the biological system, albeit with many limitations, and a model of primate color vision which is designed to self-organize in a manner that reproduces physiological data, although employing specific techniques to ensure statistical balance in the feature dimensions (Chapters 3 and 4).

In the end, we are going to discuss our results, comparing them to available psychological data collected in the past decades (Chapter 5).

Chapter 2

BACKGROUND

The relevant background to the present work can be broadly divided into three main parts. The first one deals with color vision in primates and mammals, the second tackles our current understandying of the McCollough Effect together with a complete summary of the most popular models and theoretical explanations of the phenomenon, and the last one discusses the relationship between our dissertation and previous, similar, work.

2.1 Color Vision

Most animals possess the capacity to perceive the spectral composition of light, thus experiencing color in some way. Usually, the light that is emitted by sources or is reflected by objects is composed by different wavelengths, and its specific spectrum determines the perceived color.

Different animals have different color vision capacities, with some species being able to perceive broader spectra (e.g., in the ultraviolet band) or narrower ones. Within the scope of this work we are specifically interested in mammal vision and in particular in primates'. Most mammals are generally dichromats (i.e., their retinas are sensitive to one of two wavelenth bands only, and can thus only perceive a limited color space) or trichromats, like humans (i.e., the retina is sensitive to three wavelength bands, and allows for perception of the full range of colors we are used to) [10, 26]. It is also known that some individuals (more generally women, as the genes involved are located on the X chromosome) possess tetrachromatic vision [24].

2.1.1 Retina

Visual processing in primates starts with light flowing into the eyes and onto the Retina (see Fig. 2.1). There, it is detected by the photo-receptors, special sensory cells filled with molecules that undergo a specific chain of chemical reactions when hit by photons. Photo-receptors are divided into two types, rods and cones. Rods are broadly selective to light of any wavelength in the visible spectrum, while cones have a narrow preference for determined ones. Most mammals have either dichromatic or trichromatic vision. In the case of trichromatic vision the wavelength selectivity of each cone is shaped as shown in Fig. 2.2, with peak preference around $420 - 440nm$ for the S "blue" cone (Short wavelength), $534 - 555nm$ for the M "green" one (Medium) and $564 - 580nm$ for the L "red" one (Long). Notably, despite the common naming of the three cones as Red/Green/Blue, the Long and Medium wavelength cones share a strongly correlated, similar response.

The information that is passed through the optic nerve is the further processed result of computation by Retinal Ganglion Cells (RGCs). Such cells pool different photo-receptors and filter them through receptive fields that are either ON-Center or OFF-Center (see Fig. 2.3): ON-Center cells are active when bright stimuli surrounded by darker backgrounds are seen within their receptive field, while OFF-Center cells are more active when the surround of their receptive field is brighter than the center. A similar type of ON/OFF cells are found that encode single color-opponent information: instead of being selective to contrast in luminance, they are selective to constrast in color (Long-Medium, Medium-Long, Short-(Medium+Long) [23, 9, 26].

2.1.2 Lateral Geniculate Nucleus

The Optic Nerves cross at the Optic Chiasm and finally reach corresponding structures in the opposite hemispheres, the Lateral Geniculate Nuclei (LGN). The response properties of LGN neurons are similar to those of the corresponding RGCs, with the addition of temporal lags and refinement of the receptive fields (possibly to reduce redundancy) [9]. LGN neurons project to a number of other structures. Most of their connections target the Primary Visual Cortex (V1) at different layers, keeping separate pathways segrated (magnocellular, parvocellular and koniocellular). Some of the LGN outputs reach higher cortices directly, too, bypassing V1 (e.g., V2 and V3), as found with studies related to blindsight.

The LGN of humans is composed of six layers, the first two of which are part of the

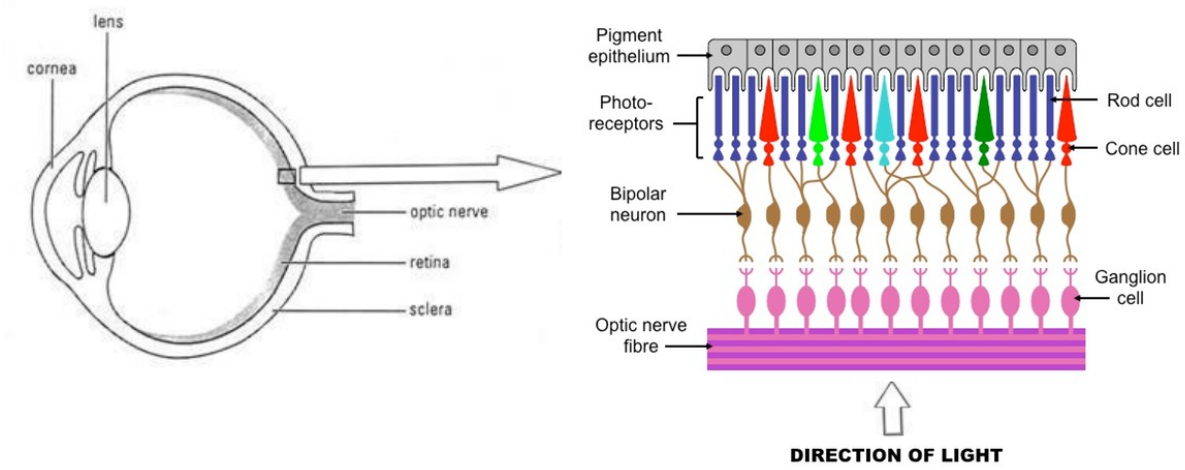

Figure 2.1: *On the left is a schematic of a primate's eye: light comes in through the cornea and the lens, and arrives onto the retina, which is shown on the right. Light is detected by the photo-receptors (rods and cones) and information is transmitted over the optic nerve. Image adapted from Scholarpedia (http://www.scholarpedia.org/article/Retina) and [http://www.ib.bioninja.com.au/options/option-e-neurobiology-and-2/e2-perception-of-stimuli.html]*

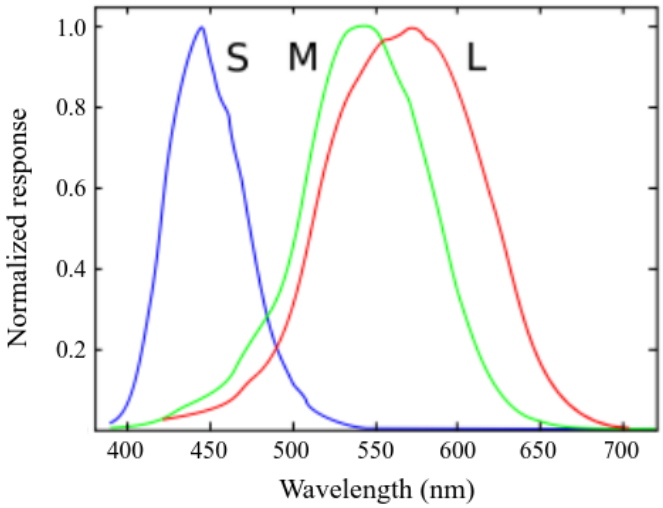

Figure 2.2: *Normalized response of retinal cone receptors in primates to light at different wavelengths (x-axis): Red corresponds to Long-wavelength cones, Green to Medium and Blue to Short. Image adapted from Wikipedia [http://en.wikipedia.org/wiki/Cone_cell].*

Magnocellular system (one for each eye) which is involved in the perception of movement and depth, while the other four are part of the Parvocellular system (two for each eye)

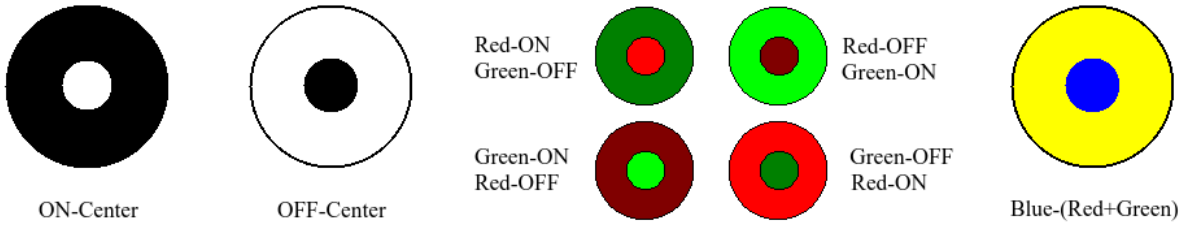

Figure 2.3: *ON-Center and OFF-Center (left) Retinal Ganglion Cells' and LGN neurons' receptive fields. Single color-opponent RGC's and LGN's receptive fields (center and right): ON- and OFF- Center cells selective to Long vs Medium (Red-Green, and vice versa) opponency and Short (Blue) vs both the other receptors.*

which is tuned for color and fine form perception. In between the M and P layers are Koniocellular cells, which are part of the system that processes information from the Short-wavelength "blue" cones.

2.1.3 Cortical Processing

Within the scope of our work, we will only model color processing at the level of the Primary Visual Cortex (V1). Further stages of processing, in particular V2 and V4, which evidence suggests are involved in the perception of hue [12] will not be discussed. We decide to follow this approach for three main reasons. First, we are interested in investigating the role of early cortical processing in color vision and in producing the McCollough Effect; second, simpler models can provide more insight into the basics of color processing in the brain, whereas it can be more difficult to get clear answers out of more complex and detailed ones; last, though the neural locus of the effect is still debated, there is good evidence that suggests that V1 could account for most of the effect.

It is known that some neurons in V1 are selective to color, and evidence was found to support the idea that color selective neurons closely match the organization of CO (Cytochrome Oxidase) blobs in the Primary Visual Cortex [29], though the response properties of such neurons are still not fully understood. Experiments found that most of the non-opponent color neurons (i.e., neurons that were found to respond to stimulation of a single type of cones, and not to the others) in V1 were also tuned to orientation (around 60% of the total color selective neurons), while single- and double-opponent neurons (i.e., neurons that receive inputs from two or more types of cones, and in the case of double-opponent, that are also sensitive to spatial contrast) were generally not [25].

Other studies however found a stricter independence between such neurons' responses to color and orientation [29].

2.2 McCollough Effect

The McCollough Effect is a visual aftereffect that links the perception of color and form (orientation) (See Fig. 1.1).

The effect was first studied by Celeste McCollough in 1965 [31] but, despite the numerous studies on the matter, we still don't understand the phenomenon fully, nor we know where is the exact neural locus where the effect is triggered.

Many studies have been done exploring the McCollough Effect and its properties. The first thing we should acknowledge is that it is long-lasting: induction times of 5 to 15 minutes can lead to a persistent effect (which continuously dacays during normal vision, but not while sleeping [30]) that can last from hours to months [27, 38]. Experiments have shown that the effect can be produced using any pair of colored gratings, but it is stronger when the colors are complementary (red and green, in particular, have been found to induce the strongest effect). Induction is possible with a single colored grating, too, in which case a form of indirect aftereffect can be observed in gratings that are orthogonal to the induction pattern. The effect is perceived the best when the spatial frequency of the gratings is similar between the induction and test stimuli. Moreover, induction patterns need not to be orthogonal, though different angles reduce the strength of the effect, up to a complete cancellation with patterns rotated less than around 11^o [17]. Also, different kinds of McCollough Effects have been found that link different types of visual modalities, like motion and color, spatial frequency and color [14], and others.

Another peculiar property of the effect is its dependency on two separate properties of visual perception, the orientation of lines and their color. Ellis [16] in particular investigated this aspect of the phenomenon, measuring the *colorimetric purity* (the amount of complementary colored light to be added to a stimulus to make it be perceived as white) of gratings at different orientations before and after induction with ME patterns : the results found that the strength of the effect, measured as the difference in colorimetric purity before and after induction, decreases linearly with the angle the gratings are rotated by, and is null at around 45^o. (see Fig. 2.4 for the plot).

Last, it is still highly debated where is the actual neural locus that is responsible for the effect. The most agreed candidate is the Primary Visual Cortex, for a number of reasons. First, the effect depends on the orientation of gratings in retinal coordinates, as

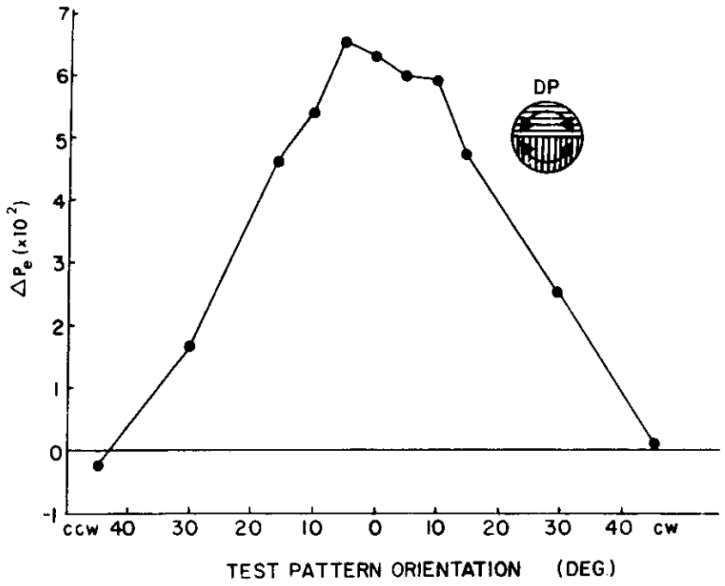

Figure 2.4: *Orientation dependency of the McCollough Effect: plot of the difference in colorimetric purity of test gratings before and after induction with colored stimuli, computed using test gratings at different orientations. Figure adapted from Ellis, 1977 [16].*

tilting a subject's head results in a weaker perception of the effect (after induction). Also, V1 is the first stage of visual processing in primates where orientation selective neurons are found, which makes it the candidate lowest in the hierarchy that could account for the effect. Moreover, no intra-ocular transfer of the effect was found [31, 18], while areas higher than V1 are know to have a high degree of binocularity.

Recent arguments, however, have been casted in favour of higher areas, first of all V2 and V4 [34], and even suggested the involvement of feedback connections. It is also been claimed that V1 is not directly linked to conscious perception [13], and could not thus be the locus of "read-out" of the effect, but critics say that conscious perception is not required at all, as it has been shown that the ME can be induced using patterns alternating at up to $50Hz$, too quick for their color differences to be perceived consciously [44]. Other evidence in support of V1 is that the McCollough Effect has been shown to respond to the wavelength of the inducing patterns rather than perceived color [42, 47], which is characteristic of V1 rather than V2, where most color-selective neurons respond to hue. Also, higher cortical lesions were found not to affect the ME [22].

In conclusion, while everybody agrees that V1 is necessary to the effect, not all agree on it being sufficient. Within our and previous work [11] we found that models of V1 can indeed reproduce the effect, though it is still possible that higher cortical areas might

have some effect under different conditions, and further investigation is necessary.

2.3 Models of the McCollough Effect

Since the first studies on the McCollough Effect [31] many different theoretical explanations have been proposed. The first idea, by Celeste McCollough herself, postulated the existence of *color sensitive edge neurons* [31, 32] that would adapt to the inducing patterns by means of homeostatic mechanisms. Any such theory that is based on homeostais requires the presence of units tuned to both orientation and color, which have been found to exist, although in a relatively small number of color selective neurons. This explanation was further suggested not to be sufficient to fully explain the phenomenon, mostly because of its extremely long lasting behavior [38], which cannot be explained by the short-term dynamics of homeostasis. It is noteworthy to say, however, that recent studies found evidence that the effect might be operating at two separate time-scales [43], and that thus the homeostatic mechanisms could account for part of the effect.

Another popular theory takes into account a purely functional description of the phenomenon, seeing it as the effect of an underlying *error-correcting system* that keeps a representation of the natural world and its long-term statistical correlations [14, 8]. Orientation and color are generally uncorrelated in the world, but during induction of the McCollough Effect they become correlated, and the error-correcting mechanisms would "correct" them. The main limitation of the model is that no neural details have been included in the model, making it somehow impractical to evaluate experimentally.

The last big theory of the McCollough Effect is based on the idea of *classical conditioning* and Hebbian learning (associative learning) [1, 21]. The theory states that the neocortex learns joint statistics of orientation and color, which can then be subjected to homeostatic mechanisms or stronger mutual inhibition. The main problem of the model is that the McCollough Effect can only be generated using specific pairs of features, like color and orientation or motion [49], and thus an approach based only on pure associative learning would not fit the effect correctly, as it would require a suitable neural representation of the features before being able to learn associations between them.

To correct the major limitations of the associative model Barlow proposed direct neural mechanisms that would implement such a behavior [3, 4, 8]. Within his model, mutual inhibitory connections are learnt by anti-Hebbian learning to decorrelate visual features, like color and orientation. During induction with colored patterns, then, the mutual inhibition between, say, vertical and red, would increase, which in turn would

inhibit the population of red selective neurons when viewing colorless vertical grids, making them appear more green. Notably a similar theory was tested successfully on the Tilt-Aftereffect [7] and the McCollough Effect itself [11]. Barlow's model is particularly interesting as it could also explain the long time-course of the effect and it doesn't require the existence of double-duty detectors.

Last, a number of neural network models were used to investigate the effect, with the most famous being the FACADE (Form-And-color-And-DEpth) general model of vision [20] and models that make use of ICA [2]. However, FACADE explains the effect as part of a complex set of mechanisms for high-level visual processing, and within the present work we are mostly interested in modelling low-level mechanisms. ICA-based models, instead, are ruled out despite being good at modelling how color and form are related in images because of their abstract relation to neural circuits, whose mechanistic explanation is the focus of this thesis.

2.4 Related Work

Apart from the models and theories we discussed in section 2.3, the most relevant previous work is the MSc thesis of Julien Ciroux [11].

Our work is extensively based on his, with a number of differences. First, we use newer and state-of-the-art models of cortical function (GCAL/AL [39], as opposed to the old one the author used (LISSOM [33]). Moreover, we work with models of color vision that are more biologically realistic in either, depending on the specific model, the functional anatomy of the system or its final cortical organization (after learning). The model used in that dissertation was related to previous works on color perception [6, 36] which inspired the present work, too, but it only modelled dichromatic vision, as it is particularly suitabile for modelling the McCollough Effect, which is (generally) induced using color pairs.

Within our simulations we used a different dataset (Barcelona Natural Images [35]) that is more balanced than the Botanical dataset used in the previous work. For measuring the strength of the McCollough Effect we used a more advanced and complete metric, too, which is necessary in order to measure this quantity in a trichromatic color space (i.e., dichromatic color spaces are one-dimensional, and measuring the colorfulness of a stimulus corresponds to computing a value on this line; trichromatic spaces, on the contrary, are more complex and different schema have to be designed in order to compute a single value for the strength.)

Finally, within our work we expand the original proposition and scientific question about the involvement of lateral inhibitory connections in the phenomenon to a more complete evaluation including homeostatic mechanisms, which where not included in the previous cortical model LISSOM [33]. As we discussed, the McCollough Effect has been found to operate on two separate timescales [43], and is thus likely that many mechanisms are involved. We investigate the role of homeostatic mechanisms in the short-term dynamics of the effect and the role of synaptic learning in lateral inhibitory connections for the longer-term ones.

In any case, we will always point to the important similarities and differences between this and any previous work within the next chapters.

Chapter 3

MATERIALS AND METHODS

The core of our project is based on the simulation of three models of color-sensitive visual systems. The underlying implementation makes use of state-of-the-art models of cortical organization. The specific models have been implemented using the Tropographica software [5] and Python scripts.

3.1 Retina and LGN

In order to describe the dynamics of Retinal Ganglion Cells (RGCs) and LGN neurons we first take input images from a color images dataset. Unless otherwise specified, the dataset we used for our simulations is the Barcelona Natural Images dataset [35]. Images were either converted to the LMS format for greater match with the biological cone-receptors or, for most of our simulations, they were kept in RGB format. Despite being less biological realistic, RGB cones have the advantage of being more decorrelated than LMS (where indeed the Long and Medium receptors have almost identical responses). This is motivated by the fact that strong decorrelation mechanisms exist at subcortical levels of color processing that further differentiate such cones' responses, and thus using RGB we can avoid modelling such specific mechanisms, still providing a good approximate behavior as it would look like at the Primary Visual Cortex.

We further model the LGN as a set of populations of neurons, according to their tuning preferences (e.g., ON- and OFF- Center color opponent units). To compute the output of each unit, each afferent (specific to the type of unit) is convolved with a Gaussian kernel, and the weighted sum over all the afferents is computed (the standard deviation of the "surround" receptive field was set to be four times the size of the "center" one). In practice the model LGN computes a Difference-Of-Gaussians response over different inputs, e.g.,

Red ON-Center and Green OFF-Surround, or Blue ON-Center and (Red+Green) OFF-Surround.

That is, the output of the LGN unit j, η_j^{LGN} is

$$\eta_j^{LGN}(t+\delta t) = f\left(\sum_{i \in F_{j,p}} \Psi_{i,p}(t) w_{ij,p}\right)$$

where $f(x)$ is a half-rectifying function ensuring that activations are positive, $F_{j,p}$ is the connection field of LGN unit j on afferent sheet p (i.e., R, G and B inputs), $\Psi_{i,p}(t)$ is the activation of the afferent unit i in such field and $w_{ij,p}$ is the connection strength between the afferent and the LGN unit, which follows a Gaussian profile specific to the projection. Each projection is identified by the pair of the afferent sheet and the LGN sheet it connects to. For instance, a Red-Green LGN ON sheet would model the ON-Center population of LGN neurons responding to Red-Green contrast, that is RedOn-Center, GreenOff-Surround, and would require $w_{ij,Red}$ to be a narrow Gaussian (Center filter), and $w_{ij,Green}$ to be a broader (negated) Gaussian (Surround filter).

3.2 GCAL - Model of Cortical Processing

In order to simulate the development and dynamics of the Primary Visual Cortex in our models we use the GCAL (Gain-Controlled, Adaptive, Laterally connected) model [39], as opposed to previous work that was based on an older version of the model, LISSOM [33]. In practice we will use the AL variant of GCAL, as we avoid using gain control for modelling the parvocellular pathway in a more biologically realistic fashion [37]. We will nonetheless refer to the model as GCAL and AL interchangeably in the text, since AL can be obtained using suitable parameter settings for GCAL. All the simulations described in the present work have been implemented and run using the Topographica simulator [5].

The model belongs to the family of Self-Organizing Maps, but is endowed with lateral connectivity (short-range excitatory and long-range inhibitory) and homeostatic adaptation. The model simulates the activity of a square sheet of cortical units (each about the level of cortical columns), each of which is connected to other units either by Afferent projections (inputs from lower stages, here only the different LGN sheets) and Excitatory and Inhibitory lateral connections (inputs from units belonging to the same simulated sheet). Each projection defines connections between input and target units, each characterized by a synaptic weight. During learning the synaptic weights of the afferent and

lateral inhibitory connections are updated, while the lateral excitatory ones are kept constant at a narrow Gaussian profile. Example connectivity can be seen in Fig. 3.1, which shows one of the color models, and features the artificial Retina, LGN, V1 (the lateral connections are drawn in yellow).

The AL model is simulated in two steps. First the activation of the units is computed, and then the plastic synapses are updated using Hebbian Learning.

The activation of V1 units is computed as

$$\eta_j^{V1}(t) = f\left(\sum_p \gamma_p C_{jp}(t)\right)$$

where p are the incoming projections (i.e., the LGN sheets and the lateral connections), γ_p is their strength scaling factor and $C_{jp}(t)$ is the contribution of sheet p to the output of unit j, which is computed as

$$C_{jp}(t+\delta t) = \sum_{i \in F_{j,p}} \eta_{i,p}(t)\omega_{ij,p}$$

$\eta_{i,p}$ is the activation of unit i from the connection field of V1 unit j in the input sheet p and $w_{ij,p}$ is the strength of the synapse linking units i (on p) and j (on V1).

$f(x)$ is a half-wave rectifying function with a variable threshold θ, updated according to the units' outputs in order to model homeostatic adaptation.

The final activation of each unit is taken after a fixed number of settling steps ($tsettle = 16$) during which the activity of V1 units is updated.

After settling, the threshold and all the synaptic weights are updated. As for the threshold, an exponential running average of the activity of each unit is updated first

$$\bar{\eta}_j(t) = (1-\beta)\eta_j(t) + \beta\bar{\eta}_j(t-1)$$

where $\beta = 0.991$ is the smoothing factor. $\bar{\eta}_j(0) = \mu = 0.024$, where μ is V1 units' target activity average. The threshold is finally modified as

$$\theta(t) = \theta(t-1) + \lambda(\bar{\eta}_j(t) - \mu)$$

$\lambda = 0.01$ being the learning rate.

Finally, the synaptic weights of all the afferent connections, together with the lateral inhibitory ones, are updated according to Hebbian Learning. Each connection field is initialized at random and enveloped in a Gaussian. Each connection is updated as

$$\omega_{ij,p}(t) = \frac{\omega_{ij,p}(t-1) + \alpha_p\eta_j\eta_i}{\sum_k \omega_{kj,p}(t-1) + \alpha_p\eta_j\eta_k}$$

where α_p is the learning rate associated to the specific input connection.

3.3 Models of color vision

Within our experiments, we used three different models of primates' color visual systems. The first two are heavily inspired by previous work [6, 36], which however was based on older models of cortical processing [33] than those used in the present work [39]. Such two models differ in that they are designed to model Dichromatic and Trichromatic vision separately. The third model we use for our investigation is based on contemporary (unpublished) work by Chris Ball, at The University of Edinburgh, and contrary to the other two models it was not modelled after the known anatomy of primates' visual systems, but was rather designed to develop physiologically plausible V1 Orientation and Hue maps. These models allow for investigation of the relative importance of various features of the anatomy and the functional organization of V1 for the McCollough Effect.

All the three systems use the same model of cortical response, AL [39], and they all share a similar structure. Every system is composed by the same model Retina (Section 3.1) and V1 (3.2), and has an intermediate layer composed by LGN sheets.

Each system has a number of LGN sheets implementing specific populations of opponent neurons, and each type of opponency is present both as ON- and OFF-Center. All LGN sheets are color-related but two *Luminosity* sheets (ON- and OFF-), and each projection from LGN sheets to the model V1 is scaled by a factor which is different for color and Luminance sheets, thus allowing for tuning of the balance of the color / form representation on the cortical surface (we will refer to this balance as the *color strength*).

It is possible to change the relative strengths of the retinal sheets (i.e., cone populations) to compensate for biases in the color distributions in the datasets, but in most of our simulations (unless otherwise specified) we will work around the problem by adding a uniformly random jitter in the input images' hue [36] (i.e., we perform uniform hue rotation on the images). We know that there exist neurons selective to any hue in the Primary Visual Cortex [48], despite the natural world being very unbalanced in the distribution of color [46]. It is likely that intrisic mechanisms exists that help rebalancing the distribution of color representations in the cortex, and while we are not going to investigate or model such mechanisms in the present work, it makes it reasonable to use hue rotation. In order to further improve the balance between the color channels, the average of the mean channel values is enforced to be fixed at 0.44 (the value was determined empirically by trial-and-error).

Each sheet is characterized by density and area parameters. The *density* is the number of simulated units (RGCs, neurons, etc..) within a unit of simulated surface, while the

area parameter sets the size of the sheet (e.g., an area of 2.0 means that the simulated surface is 2×2 bigger than the unit square, and thus has four times as many units as defined by the density parameter). Within our simulations, unless otherwise specified, the density of the Retina, LGN and V1 sheets are fixed to 24, 24 and 48, respectively, with the only difference of the idealised model that features LGN sheets of density 48. The default area parameter for V1 is set to 1.3 for the dichromatic and trichromatic models (that is, the final size is 62×62 units), and 1.0 for the idealised one. This was done empirically to optimize the cortical representation of orientation and color. The final size of the Retina and LGN sheets are computed in a similar way, with the addition of extra units on the sides to prevent border-effects (i.e., every sheets' units has a Receptive Field that lies completely within its afferent sheet) [33].

All three systems share the same parameters for the projections' scaling factors and learning rates in the model V1: the afferent strength (i.e., connections from the LGN) is set to be 1.0 in the dichromatic and trichromatic models, and 1.5 in the idealised one; excitatory and inhibitory connections' strengths are set to 1.7 and 1.4, respectively, and the learning rates of the afferent and inhibitory connections are set to 0.1 and 0.3. Lateral excitatory connections are not updated. Learning rates are divided equally among the units of a projection, and thus the effective learning rate of a single unit in a projection that has N units is

$$\alpha = \frac{\text{learning\_rate}}{N}$$

3.3.1 Dichromatic Vision

Our model of dichromatic vision is inspired to the anatomical organization of color vision in primates, but it is limited to two cone populations. Such model has many advantages. It is simple to analyse and suitable for exploring the McCollough Effect, which by itself is induced using a pair of colors, and it provides a good model of foveal vision, as the fovea doesn't have any Short-wavelength receptors (and is thus L-M dichromatic.)

Within our model the two cones are mostly uncorrelated in that we use Red and Green activations computed from RGB images rather than LMS, which however can be reasonable, as we discussed in Section 3.1.

The LGN sheets in this simplified model correspond to the ON- and OFF-Center populations of Red-Green, Green-Red and Luminosity (Red+Green) opponent neurons. The color units are double opponent, and for instance a Red-Green LGN ON neuron would correspond to a biological R+G- center and R-G+ surround. The model is shown

in Fig. 3.1.

Fig. 3.2 shows Hue and Orientation maps from an example model trained with default parameters (as we discussed in the previous section). The only modification is a lower color strength such that the total strength of the four color LGN sheets is $\frac{2}{3}$ of that of the two Luminosity sheets (0.6 strength for each color channel and 1.8 for each luminance, for total strengths of 2.4 and 3.6 respectively).

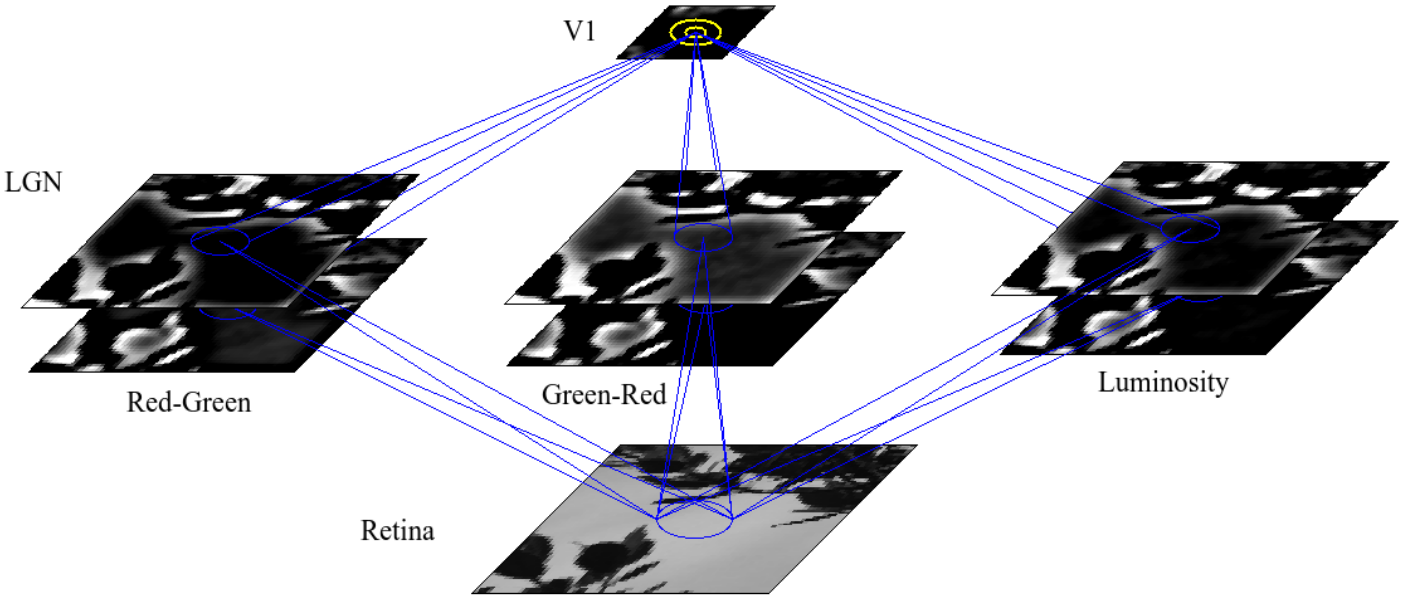

Figure 3.1: *Schematics of the dichromatic (Red-Green, Long-Medium) model of color vision. Each LGN stack represents ON- and OFF-Center populations of neurons with similar tuning (e.g., Red-Green opponency).*

3.3.2 Trichromatic Vision

The basic trichromatic vision model is similar to the dichromatic one discussed in the previous section. However, the trichromatic model further includes an anatomically inspired Short wavelength (Blue) cone receptor, together with the corresponding opponent neurons in the LGN (Blue-Center / [Red+Green]-Surround). The model's schematics is shown in Fig. 3.3, while example preference maps are shown in Fig. 3.4 (the simulation was run using the default parameters as discussed in the previous sections, and in particular the color strength was set to 0.5, thus making the whole contribution of color-related LGN sheets equal to that of the Luminance-only, colorless, channels). The dataset used for training the example maps of Fig. 3.4 is the Botanical dataset.

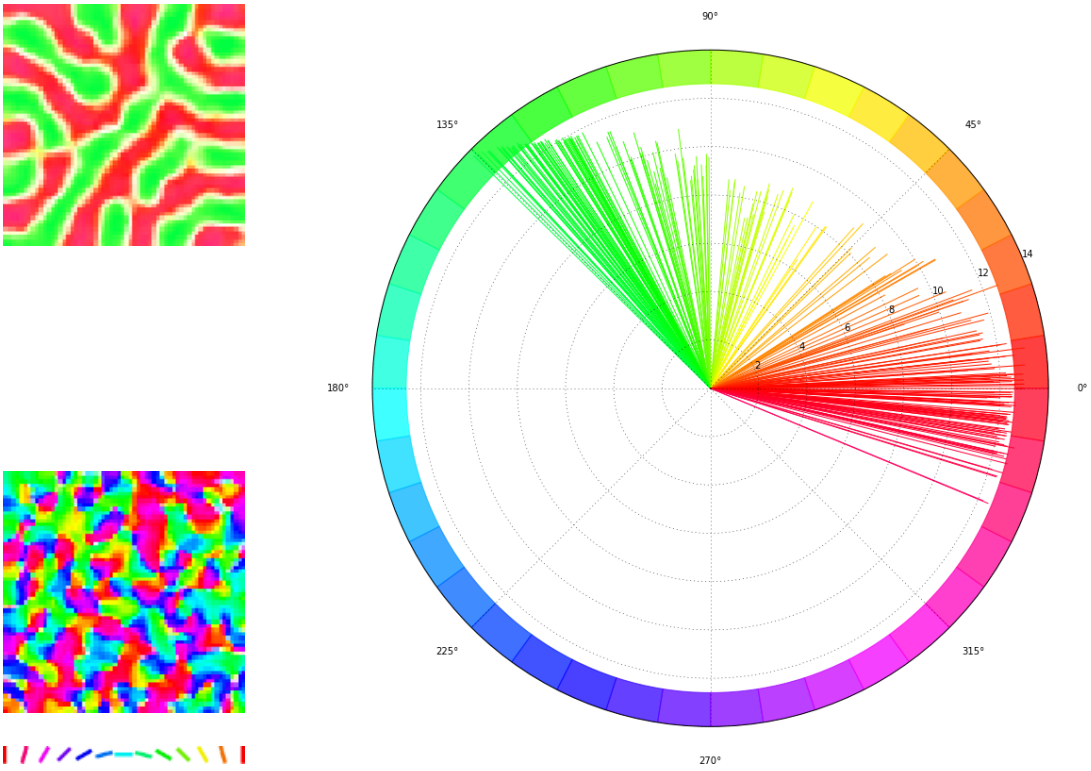

Figure 3.2: *Example cortical properties of the V1 in the* **dichromatic** *model. Left are the preference maps (top is Hue, bottom is Orientation), while Right is the polar plot of the Hue preferences, where the magnitude of the vectors represent the selectivity of each unit to the specific hue (vectors are plotted taking a reference neuron out of* 15*, to avoid crowding the plot.) The preference maps represent each unit's preference (color coded) and selectivity (color value; white pixels representing unselective neurons.)*

The model relies heavily on previous work, with crucial modifications (we use of GCAL / AL instead of LISSOM, hence including homeostatic mechanisms) [6, 36].

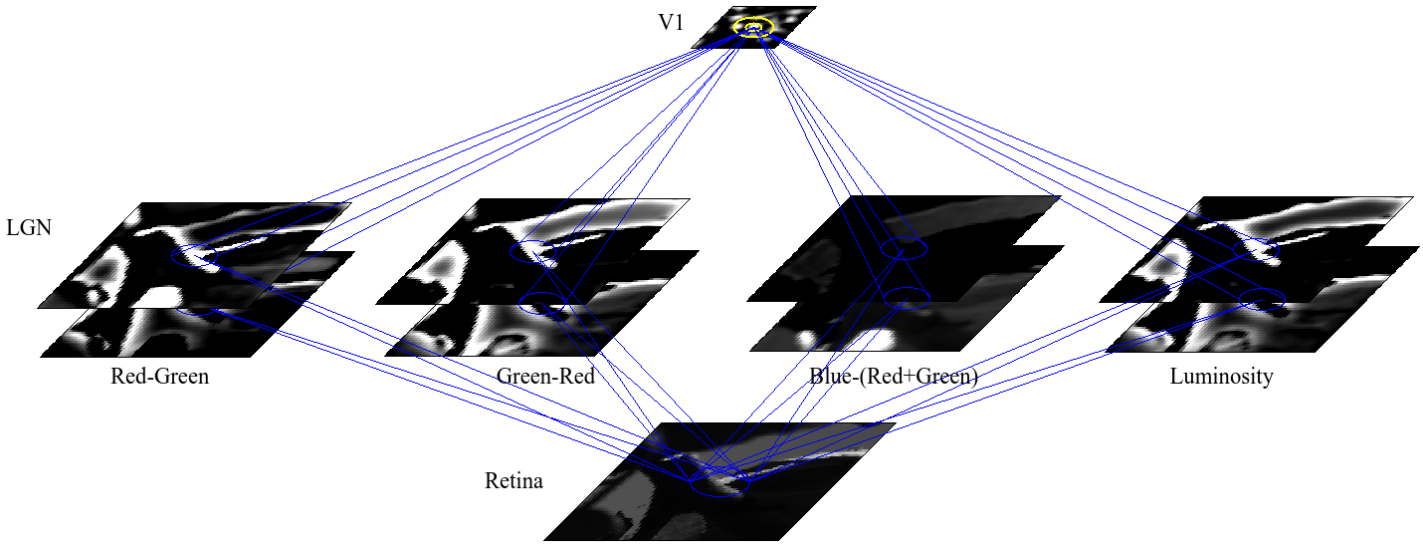

Figure 3.3: *Schematics of the trichromatic (Red-Green-Blue, Long-Medium-Short) model of color vision. Each LGN stack represents ON- and OFF-Center populations of neurons with similar tuning (e.g., Red-Green opponency).*

3.3.3 Idealised Trichromatic Vision

The final model is based on a still unpublished model by Chris Ball, developed during his PhD at The University of Edinburgh. The schematics of the model is shown in Fig. 3.5, while example preference maps are reported in Fig. 3.6.

The model is named *Idealised* because its objective is to provide a (not necessarily biological) subcortical system such that the resulting cortical maps in V1 would be as close as possible to the physiological data available. As can be seen in Fig. 3.5, the LGN is simulated to have populations of every color opponency (e.g., Red-Blue, which is not found in any trichromatic animal), achieving a perfect symmetry between the different cones. Moreover, the total activity of the LGN sheets is re-normalized such that at any iteration the mean activity of the sheets is set to a fixed value, here 0.05.

3.3.4 Training

Training was performed by extracting random patches from the Barcelona natural images dataset [35] and presenting them to the models.

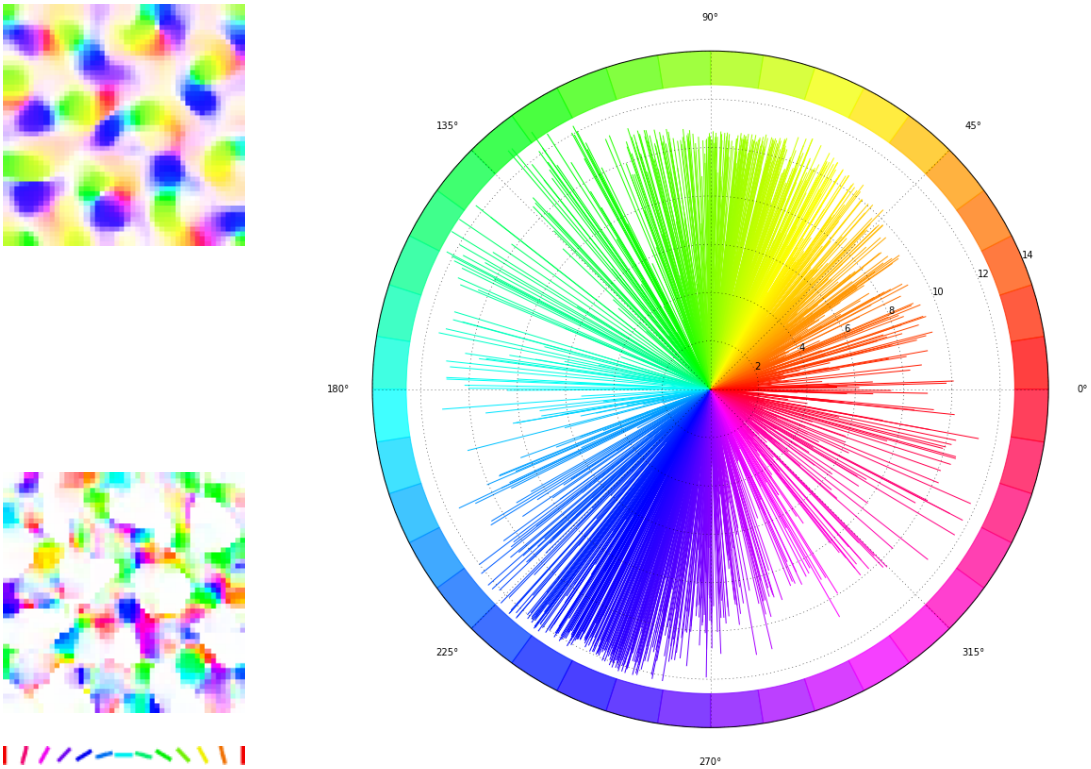

Figure 3.4: *Example cortical properties of the V1 in the* **trichromatic** *model. Left are the preference maps (top is Hue, bottom is Orientation), while Right is the polar plot of the Hue preferences, where the magnitude of the vectors represent the selectivity of each unit to the specific hue (all units are plotted.)*

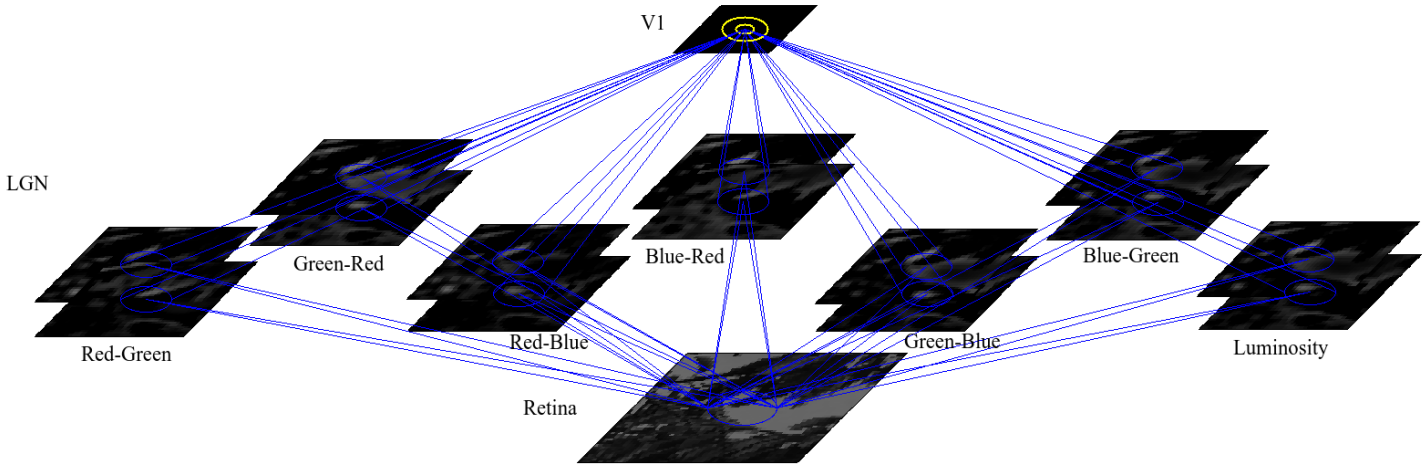

Figure 3.5: *Schematics of the idealised model of trichromatic vision. Each LGN stack represents ON- and OFF-Center populations of neurons with similar tuning (e.g., Red-Green opponency).*

It is known that cortical learning modelled using LISSOM / GCAL results in cortical maps that do not depend on the initial weights' values, but rather only depend on the input stream, that is the exact sequence of images used during training [33]. We effectively use this fact to simulate different subjects for higher statistical relevance, using different values for the random generators' seeds.

The training of the model was performed by presenting input patches and simulating the cortical response, and then updating the cortical model as described in Section 3.2. Training lasted 10000 iterations for most of our simulations, which is around the number of iterations generally performed with the model we use. Training for a longer time didn't result in significantly better results on testing the McCollough Effect, and thus the final number of iterations was decided from a computational efficiency point of view.

3.4 McCollough Effect

The specific modelling of the McCollough Effect was divided into three phases: testing before induction, induction and testing after induction. The final results compare the decoded colorfulness of stimuli before and after adaptation to the colored gratings.

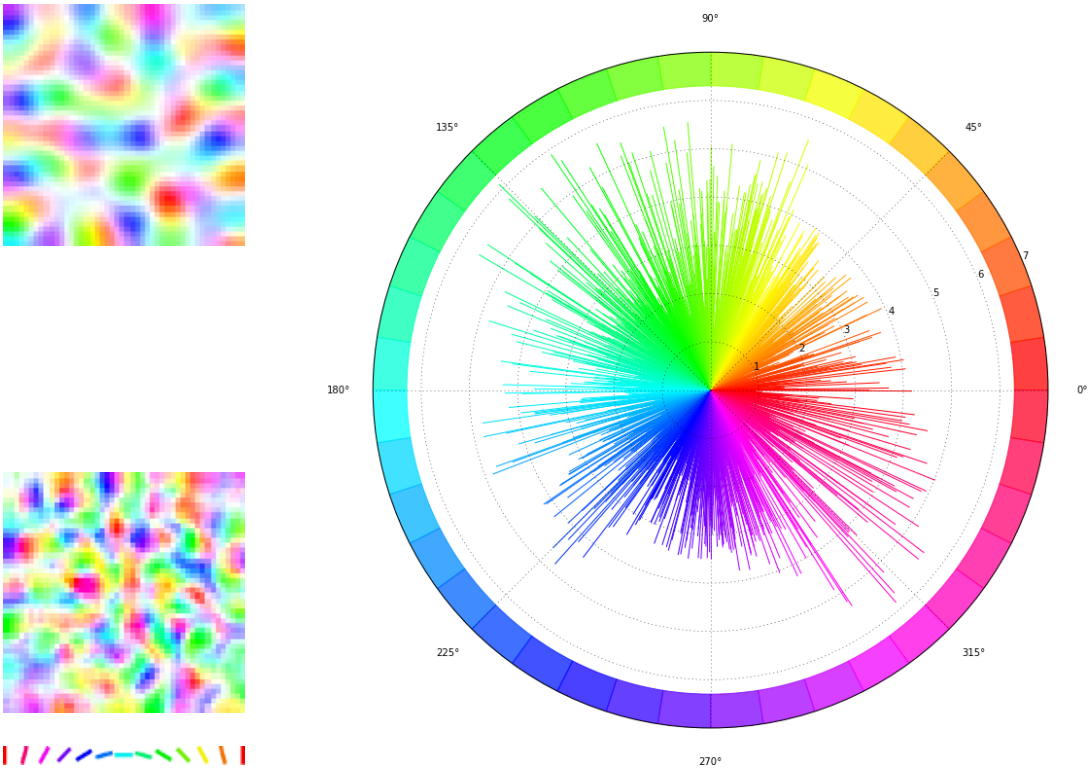

Figure 3.6: *Example cortical properties of the V1 in the* **idealised trichromatic** *model. Left are the preference maps (top is Hue, bottom is Orientation), while Right is the polar plot of the Hue preferences, where the magnitude of the vectors represent the selectivity of each unit to the specific hue (all units are plotted.)*

3.4.1 Induction and Testing

The induction of the McCollough Effect was based on previous work [11]. The patterns used in the present work are colored sine gratings, which are set by drawing gratings of different amplitudes on the different retinal sheets, thus simulating the activation of the different cones. During stimulation with the induction patterns the afferent connections (LGN to V1) were set to have a null learning rate, thus only allowing lateral inhibitory connections and homeostatic mechanisms to adjust themselves. In previous work, aftereffects were found to be dominated by the plasticity of lateral connections [7], and here we can test such hypothesis on the McCollough Effect by separately enabling and disabling cortical plasticity and homeostatic mechanisms.

Each of the inducting pattern was presented in alternation for 5 iterations each time, repeating the sequence until a total of 300 presentation steps were simulated for each pattern. The simulated length of the induction was found to be in line with psychological experiments on human subjects, as previous work used a correspondence of 90 iterations for three minutes of induction (Tilt Aftereffect, [7]), thus making 600 iterations suitable for modelling a $15 - 20$ minutes McCollough Effect test.

The phase of the induction gratings used in each 5-iterations block was sampled from a normal distribution with a variance of $\frac{1}{8}$ of the gratings' period, effectively modelling a fixed gaze experiment.

3.4.2 Decoding the Perceived color

The problem of decoding the color that is perceived by a cortical model at a given time is difficult for many reasons. We know that specific visual properties can be decoded statistically using population codes in V1 [19], but problems might rise if the cortical model is not well balanced (i.e., some features are over-represented). Also, color is not unidimensional: in the case of dichromatic simulations any perceived color can lie only on a line, that goes from "more color A" to "more color B". In a general color system however this is not true anymore.

In order to tackle this problem, we decided to take advantage of the simplified context of dichromatic simulations, which are a good fit for studying the McCollough Effect, in that color pairs are used. We can thus compute a measure of the colorfulness of a pattern along an axis where the two opposite directions encode for the perception of two different colors. In practice, we assign a color vector to each V1 unit, whose argument is the neuron's preferred hue and its magnitude is be the product of the neuron's selectivity to

that hue and its activation. We proceed by computing a unit vector $\vec{v}$ whose argument is that of the average vector over the population,

$$\angle\vec{v} = \arctan \frac{\sum_i \phi_i \eta_i \sin(h_i)}{\sum_i \phi_i \eta_i \cos(h_i)}$$

where η_i is the activation of unit i, h_i is its hue preference ofand ϕ_i its selectivity.

Then, we compute the dot product between the vector $\vec{v}$ and the unit vectors representing the two colors we want to compare (e.g., Red and Green), and compute the difference (the sign is arbitrary, depending on which color we take as positive), but within this work we'll take Red to be positive. See Fig. 3.7.

The result is a single number which is close to zero (provided the neural population has balanced hue preferences) when colorless stimuli are presented as input, and increasingly positive (negative) for stimuli with the first color of the pair (the second). In the case of Red-Green comparison, within the scope of the McCollough Effect, stimuli that are perceived as "Red" will have a high positive value, while stimuli perceived as "Green" will have a high negative value. It is also worth noting that any baseline dependent on the balance of neurons' hue preferences can be evened out by computing differences between such measures of colorfulness (e.g., before and after induction of the McCollough Effect).

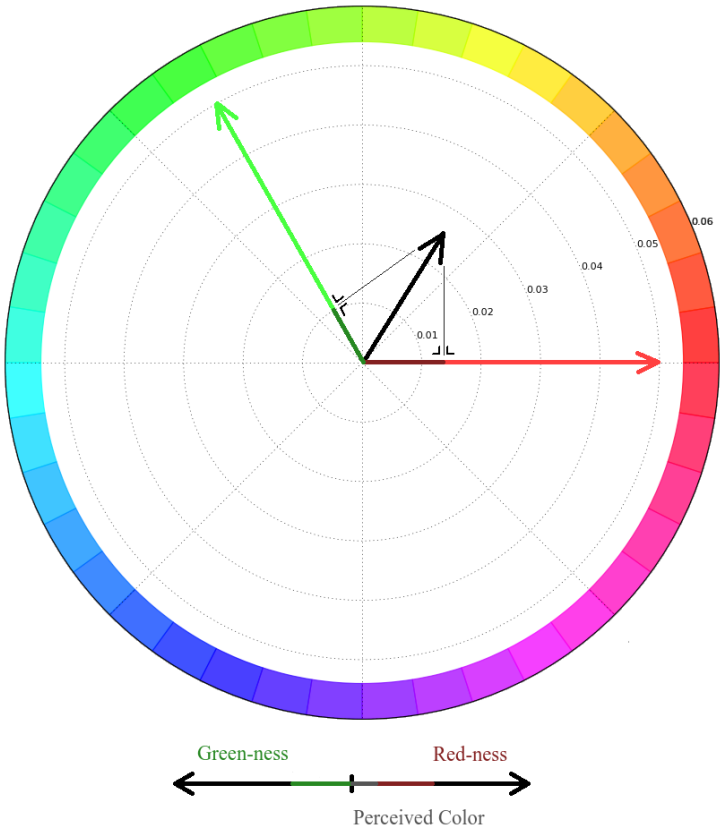

Figure 3.7: *Example decoding of an average population vector scaled by hue preferences and selectivities of V1 neurons (black). Red and green are reference vectors onto which the test vector is projected (dark red and dark green).*

Chapter 4

EXPERIMENTS AND RESULTS

Within this chapter we are going to present different aspects of the McCollough Effect for which physiological data is available for comparison. We will try to compare our different models wherever possible and we will prepare the basis for further discussion about the models and the dynamics of the McCollough Effect, which we will explore further in Chapter 5.

4.1 The McCollough Effect

We start with the most basic experiment to discuss the first properties of the phenomenon. In this setting, we present two colorless gratings with orthogonal orientation and same spatial frequency (in particular, one vertical and one horizontal) before and after induction with colored gratings, and we compute the difference in colorfulness before and after adaptation. We will present the exact values of the colorfulness of the white gratings before and after induction, too. However, as there are biases dependent on the unbalance of hue-preferences in the model V1, their difference will be the most meaningful measure.

Induction is performed using similarly oriented gratings, under two different experiments for each model: Red-Vertical, Green-Horizontal for the first test and Green-Vertical, Red-Horizontal for the second one. In both tests we expect the perceived color of each test grating to be shifted towards the color opposite to the corresponding induction pattern.

For the experiments with the **dichromatic** model we trained 10 models using different random seeds to train them using different input streams, thus effectively simulating different subjects' visual experiences and achieving different neural layouts. Each model was trained for $N = 10000$ iterations, and slightly more strength was given to Luminance

LGN sheets compared to color-specific ones ($color\_strength = 0.4$) in order to improve the orientation map in the model V1. Random hue-jitter was also applied to input images in order to obtain a more balanced representation of colors in the inputs, independent from the specific natural images statistics (e.g., abundance of green and blue, for woods and the sky). We computed the required measures for each model and then their average and standard deviation. The results are shown in Table 4.1. Fig. 4.1 shows example population activities and average population vectors. As we can see, the perception of colorless gratings is shifted towards the color of the other induction grating. Vertical gratings will thus look greener" when Red-Vertical patterns are used for induction, and "redder" when Green-Vertical patterns are used. We also observe a good balance between the shifts in the two directions (Red and Green).

TEST 1	Before	After	Difference
VERTICAL	-0.154779 (0.176994)	-0.626384 (0.191359)	**-0.471605** (0.139545)
HORIZONTAL	-0.042584 (0.174001)	0.438041 (0.150593)	**0.480625** (0.137430)

TEST 2	Before	After	Difference
VERTICAL	-0.154779 (0.176994)	0.396533 (0.191535)	**0.551312** (0.145786)
HORIZONTAL	-0.042584 (0.174001)	-0.555774 (0.157620)	**-0.513190** (0.155681)

Table 4.1: **Dichromatic** *system. Measured colorfulness of white gratings Before and After induction with colored gratings. The values represented the perceived color, as Green (below zero) and Red (above zero). The data shown is the average over* 10 *simulated subjects, and it's reported as mean values with their standard deviation (in parentheses). Top table: first test, induction with Red-Vertical and Green-Horizontal gratings. After induction colorless vertical gratings appear greener (negative) and horizontal ones appear redder (positive). Bottom table: second test, using Green-Vertical and Red-Horizontal; vertical gratings look redder (positive) while horizontal ones look greener (negative).*

For the **trichromatic** model we simulated 5 different subjects in the same way as we did in the previous experiment. The color and the luminance LGN sheets are given a total equal weight (i.e., the two luminance sheets have a strength that is equal to that of all the color sheets). We computed the strength of the McCollough Effect for each subject, and we extracted statistical information, which is reported in Table 4.2. As we

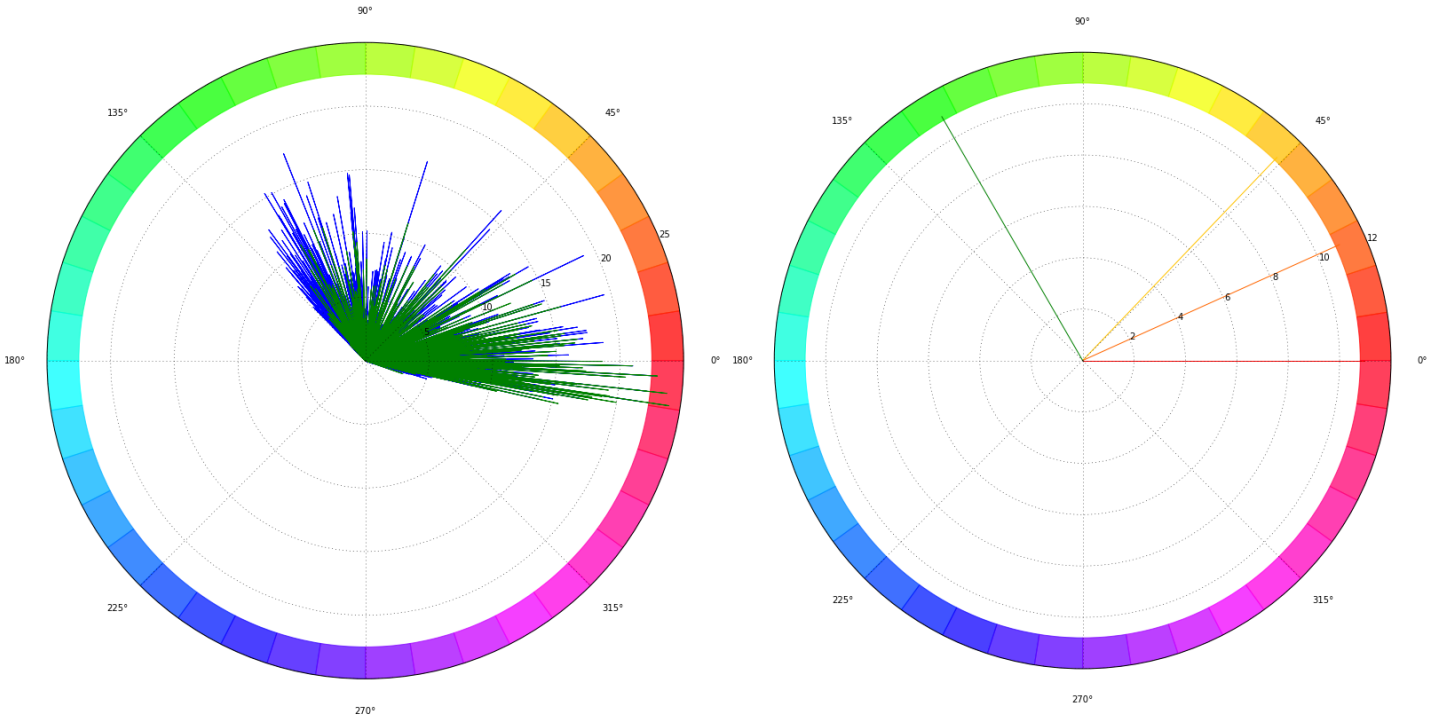

Figure 4.1: *Example population activity before and after induction in the* **dichromatic** *model. Left is the population activity drawn using hue-preferences as the vectors' arguments and their selectivity times activation as their magnitude. Blue vectors are the responses to a white horizontal grating before induction, Green vectors are the same after induction with a Green-Horizontal grating. Right picture shows the vector averages: light orange is the average before induction, darker orange is the average after induction. Red and Green vectors are shown as reference.*

can see from the variance of the values, the system is more unstable than the previous simulation, which was well balanced and consistent. This can be for many reasons, and a better understanding of the differences between the two simulations will be interesting future work. As far as we could analyze the models, we found that the first one had neurons that represented most orientations with good accuracy (consistent with colorless previous simulations using GCAL, which had a mean decoding error of 5^o to 10^o,) while the second was poorer in such a representation. The representation of color, on the other hand, was very balanced in both models.

Finally, in the case of the **Idealised Trichromatic** model we simulated another 5 test subjects and performed induction with the appropriate gratings as described in Section 3.4.1. The results are reported in Table 4.3, and similarly to the case of the

TEST 1	Before	After	Difference
VERTICAL	0.028121 (0.841249)	-0.176233 (0.735258)	**-0.204353** (0.335443)
HORIZONTAL	-0.230547 (0.901317)	0.084596 (1.033288)	**0.315143** (0.210030)

TEST 2	Before	After	Difference
VERTICAL	0.028121 (0.841249)	0.767211 (0.880755)	**0.739090** (0.563988)
HORIZONTAL	-0.230547 (0.901317)	-0.770934 (0.479388)	**-0.540387** (0.432896)

Table 4.2: **Trichromatic** *system. Measured colorfulness of white gratings Before and After induction with colored gratings. The values are Green (below zero) and Red (above zero), and represent the average over* 5 *simulated subjects, reported together with their standard deviation (in parentheses). Top table: first test, induction with Red-Vertical and Green-Horizontal gratings. After induction vertical gratings appear greener (negative) and horizontal ones appear redder (positive). Bottom table: second test, using Green-Vertical and Red-Horizontal; vertical gratings look redder (positive) while horizontal ones look greener (negative).*

trichromatic model, the variance is very high, for similar reasons. Interestingly we found that the simulations that had the strongest effect had the lowest orientation decoding error on average, while simulations that failed in achieving a good representation of the orientation space could not reproduce the McCollough Effect. This is reasonable because the McCollough Effect relies on the interactions between the perceived orientation of patterns and their color, and a poor representation of either of these features is likely to prevent the effect to be perceived.

Moreover, and critical for the following sections, we could only achieve a reasonable representation of both color and orientation in the *dichromatic* model, with the other two achieving partial representations and thus partial results.

4.2 Orientation dependency of the McCollough Effect

Ellis [16] studied the dependency of the strength of the McCollough Effect on gratings with different orientations (i.e., not only vertical and horizontal). Fig. 2.4 shows her data, which is characterized by two main features: gratings oriented around 45^o off the induction gratings appear to be colorless, and the decay is the strength of the effect is

TEST 1	Before	After	Difference
VERTICAL	0.339216 (1.296216)	-0.216267 (0.231275)	**-0.555484** (1.168237)
HORIZONTAL	0.033844 (1.482212)	0.483650 (0.172860)	**0.449805** (1.390989)

TEST 2	Before	After	Difference
VERTICAL	0.339216 (1.296216)	0.467627 (0.128034)	**0.128411** (1.207255)
HORIZONTAL	0.033844 (1.482212)	-0.325781 (0.219666)	**-0.359625** (1.413137)

Table 4.3: **Idealised Trichromatic** *system. Measured colorfulness of white gratings Before and After induction with colored gratings. The values are Green (below zero) and Red (above zero), and represent the average over* 5 *simulated subjects, reported together with their standard deviation (in parentheses). Top table: first test, induction with Red-Vertical and Green-Horizontal gratings. After induction vertical gratings appear greener (negative) and horizontal ones appear redder (positive). Bottom table: second test, using Green-Vertical and Red-Horizontal; vertical gratings look redder (positive) while horizontal ones look greener (negative).*

linear to the angular divergence of the test patterns with comparison to the induction ones.

In the case of our **dichromatic** model we replicated a similar figure by averaging over 10 simulated subjects, using the same parameters as in Section 4.1. The resulting plot is shown in Fig. 4.2. As we can see, not only the plot resembles Ellis' in great detail, but it also perfectly matches the described features.

As we discussed in Section 4.1, the overall quality of the representation of orientations in the other two models is lower, and this makes it difficult to produce a clear plot like Fig. 4.2.

4.3 Different Orientation of the Inducing Patterns

The final test we performed tried to replicate the psychological data collected by Fidell, in 1970 [17], which was also studied in the previous thesis by Ciroux [11], thus allowing for a direct comparison between our models and human data. For the scope of the present simulations we only used the **dichromatic** model because of its good representation of

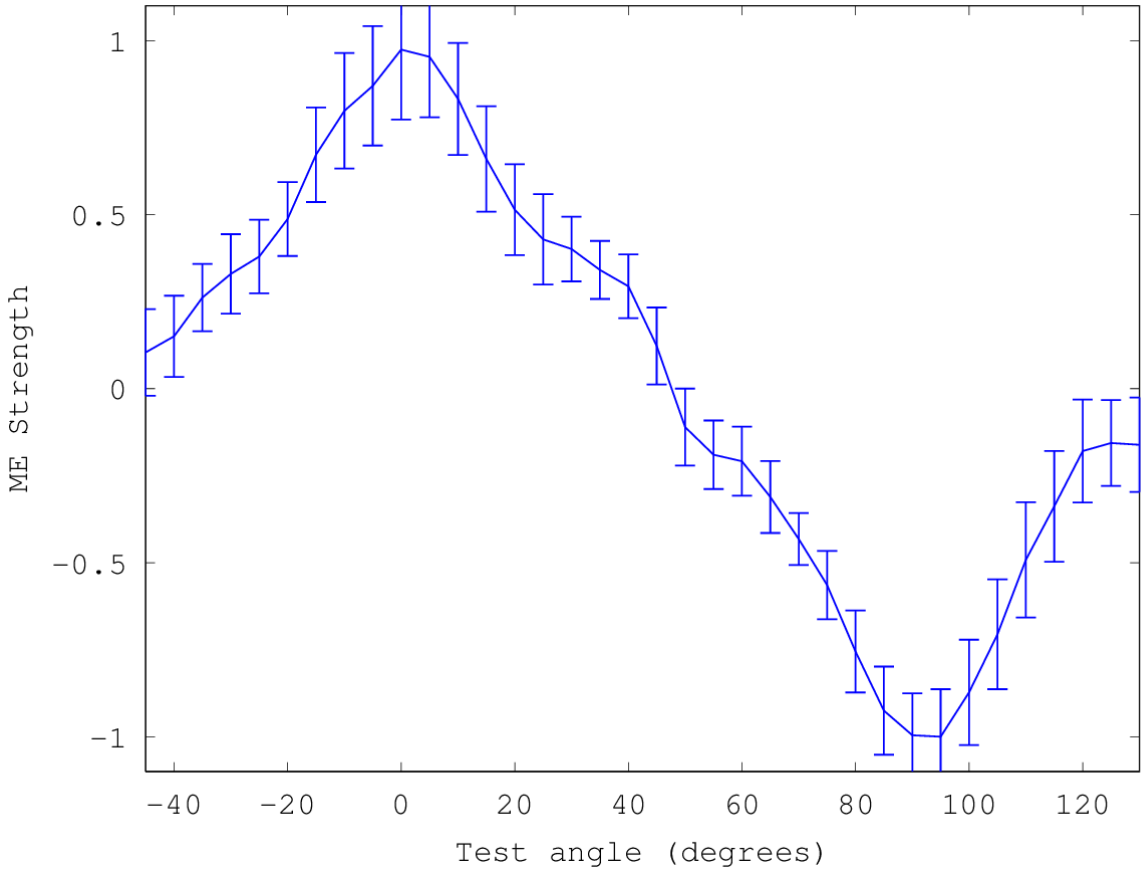

Figure 4.2: *Orientation dependency of the ME in the* **dichromatic** *model, averaged over* 10 *simulated subjects. The strength shown is the difference in colorfulness of white gratings before and after induction with Green-Horizontal (*0^o*) and Red-Vertical (*90^o*) gratings. Positive values mean that the perception is shifted towards Red. Bars represent the standard deviation of the measures.*

orientations, which was required for the experiments.

Fidell used pairs of gratings at different orientations during induction and testing (induction and testing was always performed with the same gratings, except for their color). In the first experiment (Fidell, 1, cyan line, in Fig. 4.3), as a control, induction was performed using a Red grating oriented at the angle shown on the x-axis, and a Green grating orthogonal to it. The strength of the McCollough Effect, measured using test gratings with the same orientation as the induction ones, was found to be independent from the exact orientation of the patterns, as long as they were orthogonal. Our results for the same setting are shown in Fig. 4.3 (experiment 1, blue line), with the difference that we could afford computing more sample points than in the original experiment. Each

point corresponds to the average of 5 simulated subjects, with the bars representing the standard deviations. Our simulations show the independence from the orientation of the induction patterns (as long as orthogonal) with even more detail than the original data.

In the second experiment, Fidell fixed the Green grating to always be vertical (90^o), while the orientation of the Red grating was set to the value on the x-axis. In this setting, the induction gratings are increasingly closer together, and it is hence possible to compute the strength of the McCollough Effect that results from induction patterns oriented at different angles (i.e., not orthogonal). Fig. 4.3 shows a comparison between Fidell's data (Fidell, 2, magenta line) and ours (experiment 2, red line), which are remarkeably similar, showing a significant decay in the strength of the effect when patterns are less than 45^o apart, rapidly falling and becoming almost non-existent when the difference in the orientation of the gratings is around 10^o. Again, the computational approach allows for sampling more points of the curve.

Our results match the previous work by Ciroux, too (see Fig. 4.6 of [Ciroux, 2006] ([11]), but in the case of our experiments we used a more accurate method for measuring the strength of the effect. For every simulation (i.e., pair of gratings at specific orientations) we presented both combinations of Red and Green with each orientation, and computed the average of the four values (sign-corrected) of the measured strength (i.e., 2 test values corresponding to each induction grating, for each association Red-first grating, Green-second grating and Green-first grating, Red-second grating). For example, in the default case of horizontal and vertical gratings we'd compute the differences in the perceived colorfulness of the test patterns in each of the two combinations (obtaining two values for each combination, one positive, perceived as "redder" and one negative, perceived as "greener"). The resulting strength is then the average of the absolute values of the 4 measurements.

4.4 Role of Homeostasis and Lateral Inhibitory Connections

One critical aspect we tried to investigate is about the specific contribution of homeostatic mechanisms and cortical plasticity (of lateral inhibitory connections) to the production of the effect.

One advantage of running computer simulations is that, contrary to the biological system, we can selectively disable either mechanism, and we can thus compare the differences

in the McCollough Effect that each of them triggers separately.

In particular, we run two sets of experiments, both using our dichromatic system. In the first one, we compared three scenarios, namely one with both mechanisms enabled (i.e., homeostatic adaptation and plastic lateral inhibitory connections), another with only homeostasis enabled, and the last one with only plasticity, and we compared the strength of the effect as measured using colorless gratings with the same orientation as the induction patterns (horizontal and vertical). The results are reported in Table 4.4, and compare the three cases ("default" meaning that both mechanisms are enabled) in the two experimental settings, Red-Vertical, Green-Horizontal (Test 1) and Green-Vertical, Red-Horizontal (Test 2), averaging each value over 10 simulated subjects. As we could expect, the effect is stronger when both mechanisms are present: this can also fit the finding that the McCollough Effect operates at two different timescales, likely resulting from two different mechanisms active at the same time [43]. In the case of the experiments run enabling a single mechanisms, the results from enabling homeostasis (and disabling cortical learning) are slightly stronger and more balanced than by only enabling cortical plasticity. Even though this could hint at a stronger dependency of the effect on homeostasis, it is unlikely to explain its long-lasting dynamics. Moreover, the test patterns in the experiments share the same orientation as the induction gratings, which would be expected to work using homeostatic mechanisms in any simulation that had double-duty neurons (that is, neurons selective for both orientation and color).

In order to further investigate the role of homeostatic adaptation and cortical learning in the McCollough Effect we run a set of simulations similar to the first ones we described, but instead we plotted the strength of the effect when using test patterns at different orientations (like in Section 4.2). Comparing the effect using data from all the orientations is useful to explain the stronger effect observed when enabling the only homeostatic adaptation. Indeed, we would expect the effect to be very strong near the orientations used for induction, because of the adaptation to such stimuli, but not to be able to account for the full effect as was measured in Section 4.2.

We show the results in Fig. 4.4, where each of the two plots compares the strength of the McCollough Effect tested at different orientations using the default model (both mechanisms enabled) and each of the two models where only one is enabled at a time. Notably, as we would expect, when only homeostasis is enabled the peak of the curve matches the default simulation: orientations similar to the induction patterns are *directly* affected by adaptation. However, it is interesting that in a such setting the width of the curve is much narrower than the curve computed using the default simulation, and in

TEST 1	Default	Only Homeostasis	Only Plasticity
VERTICAL	-0.471605 (0.139545)	-0.459404 (0.103371)	-0.342562 (0.168272)
HORIZONTAL	0.480625 (0.137430)	0.454146 (0.139887)	0.394118 (0.152049)

TEST 2	Default	Only Homeostasis	Only Plasticity
VERTICAL	0.551312 (0.145786)	0.497714 (0.147248)	0.453383 (0.216519)
HORIZONTAL	-0.513190 (0.155681)	-0.517360 (0.157267)	-0.390162 (0.198635)

Table 4.4: **Dichromatic** *system. Measured colorfulness of white gratings before and after induction with colored gratings, averaged over* 10 *simulated subjects. Positive values indicate a shift in perception towards Red while negative values are associated to Green. The standard deviation of each measure is shown in parentheses. The two tables show data from two tests, which differ in the induction patterns used: Test 1 uses Red-Vertical and Green-Horizontal gratings, while Test 2 uses Green-Vertical and Red-Horizontal ones. After induction in Test 1 vertical gratings appear "greener" (negative values) while horizontal ones appear "redder" (positive values), and the opposite is true for Test 2. Note that differently from the tables in Section 4.1, here we only reported the differences in the perceived colorfulness of the test stimuli, that is the measure of the strength of the effect, and not the exact values before and after induction.*

particular the effect decays at around 30^o, as opposed to the psychologically measured 45^o in humans (note that this is partially similar in the default simulation, for the gratings around 90^o, right part of the graph).

In contrast, when only cortical plasticity is enabled, and homeostasis is disabled, the peaks match the default curve worse (they appear to be more "flat"), but the width of the curve is a perfect fit to human data, which is indeed matched better than in the default simulation. We suggest that the closer match hints at a major involvement of cortical plasticity in the McCollough Effect, and that the discrepancy near the induction orientations is actually due to the poor representation of orientation in the model, which is on average off by 5^o to 20^o, thus introducing errors in the plot.

In conclusion, homeostasis seems to be describing the effect around the orientation of the induction gratings very well, though longer-term effects would need to be considered (i.e., at some point in time the adaptation would fade out, leaving only the effect of the

updated synapses), while cortical plasticity probably accounts for the accurate shape of the orientation-dependent curve as measured in humans (i.e., linear decay of the effect with the orientation of the test patterns and width of the curve).

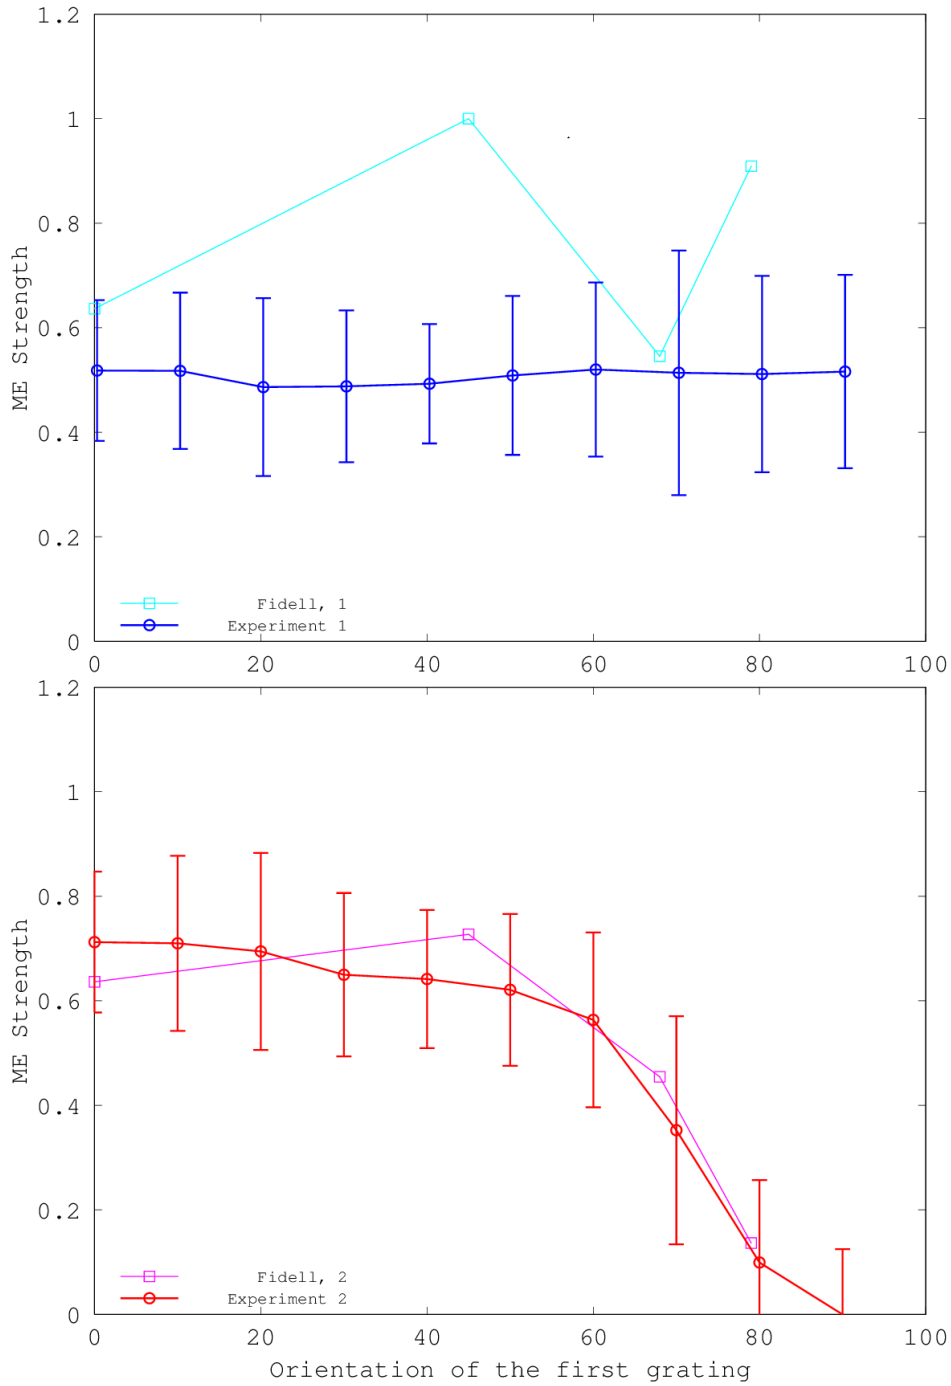

Figure 4.3: *McCollough Effect : dependency on the angle between the induction gratings. Cyan and Magenta are data from [Fidell, 1970] ([17]), while Red and Blue are the experiments from the present work. The* **top** *figure shows a comparison between our data and Fidell's first experiment. It represents the average strength of the ME inducted using orthogonal gratings, the first of which being oriented as the angles shown on the x-axis (e.g., points at* 30^o *stand for inducing patterns oriented at* 30^o *and* $30^o + 90^o = 120^o$*). The strength of the effect is constant regardless of the gratings used, as long as they are orthogonal. The* **bottom** *figure shows the strength of the ME obtained using induction gratings with different orientations: the first grating is oriented at the angle on the x-axis while the second is always vertical* (90^o)*. The strength of the effect is maximal when the induction gratings are orthogonal, and it decreases quickly for patterns that are less than* 45^o *apart. Data from our experiments is averaged over* 5 *simulated subjects.*

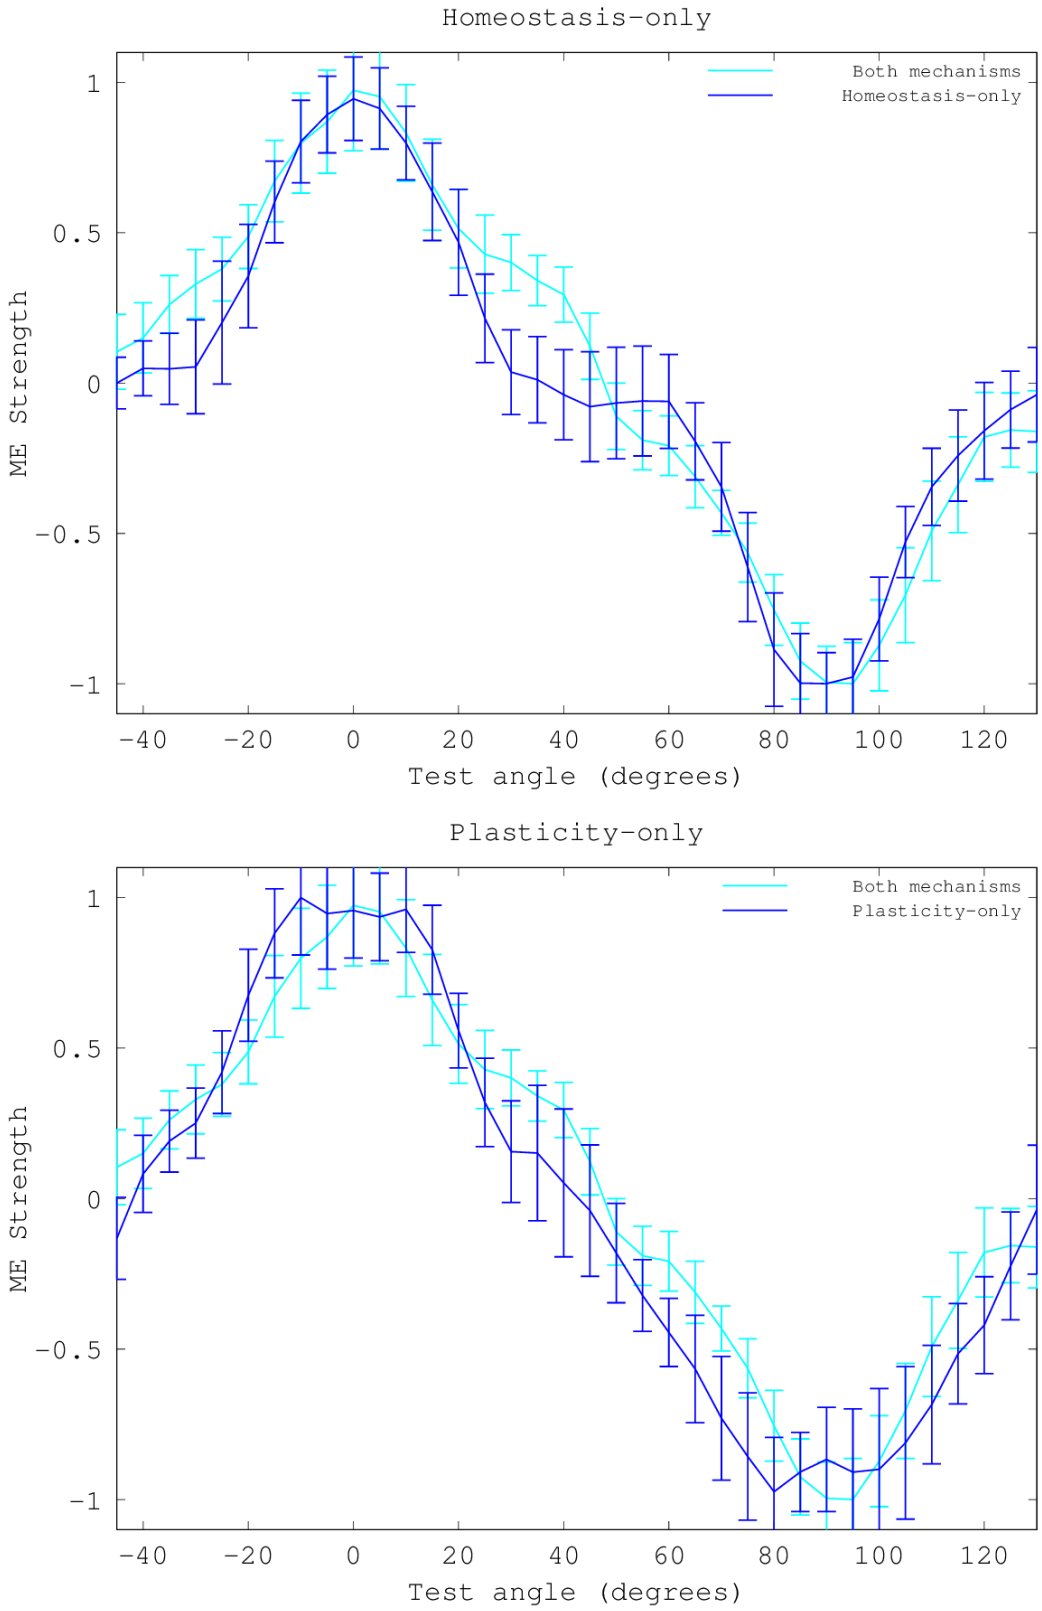

Figure 4.4: *Comparison between the dependency on the orientation of the test patterns of the McCollough Effect, here obtained using Green-Horizontal and Red-Vertical gratings. The top plot shows the strength of the McCollough Effect for gratings at each orientation when only homeostatic adaptation is enabled (and cortical plasticity is disabled). The bottom one shows similar data, obtained using the reverse combination (cortical plasticity is enabled and homeostasis is disabled). In both plots a reference curve is shown in cyan, and corresponds to the strength measured when both mechanisms are enabled.*

Chapter 5

DISCUSSION AND CONCLUSIONS

In Chapter 4 we presented data from a number of experiments that we performed using our models. In particular, we could reproduce the McCollough Effect under a variety of conditions, exploring most of its characteristics and comparing our results to available physiological data.

In the specific case of dichromatic simulations we could reproduce the results from the closest previous work [11] using the new models and GCAL, and we could compare some of the similarities and differences between the models. The results are interesting and accurate, though we found a major difference in the learnt models in the number of "double-duty" units (units that are selective to both orientation and color), which in our models were predominant. Interestingly, having more double-duty cells makes it easier for homeostatic mechanisms alone to reproduce the effect, even though in biological systems their number is lower, and so is probably the contribution of the homeostasis. Further, it is difficult to compare our model with Ciroux's [11] qualitatively as both works match the available human data perfectly[16, 17].

5.1 Summary of the results

We first reproduced the McCollough Effect using our three models, and we found it easier for the dichromatic system to produce a stronger, balanced effect than using the others. This could be due to many reasons, the most critical being that, even though all three models have learnt a very balanced representation of the input color space (with great help of the hue-rotation transformations), their representation of orientation was found to be poor (as measured by decoding the orientation of white gratings and computing the error). We tested the accuracy of the representation by presenting gratings at different

orientations and decoding the perceived orientation (using a population average decoder), and finally computing the error between the real and the decoded values. Between the three models, the dichromatic system had the best accuracy (even though the error was still considerable), while the other systems failed with even higher errors. As the effect relies on the precise interactions between the representation of color and orientation, an inaccurate coding of orientation can effectively prevent the effect from matching the psychological data.

We further tested this hypothesis while also validating our models presenting test patterns (before and after induction) at different orientations, with the scope of reproducing the data collected by Ellis in 1977 [16], which shows the dependency of the strength of the McCollough Effect on the orientation of the test gratings. As we saw in Section 4.2, the results from the dichromatic simulations matched closely the data, meaning that the model is capable of reproducing the McCollough Effect to a great level of realism. Another test we ran to determine the behavioral match between our dichromatic model and the psychologically-measured effect was based on the experiments by Fidell in 1970 [17], which tested the relationship between the mutual orientation of the induction gratings and the strength of the resulting McCollough Effect (e.g., for determining the minimum difference in orientation capable of inducing the effect). Again, our results match closely the human data and are in agreement with previous results [11].

5.1.1 Neural mechanisms underlying the McCollough Effect

Having determined the level of behavioral match between our model and human data from previous works we investigated the possible neural mechanisms capable of producing the McCollough Effect. In particular, we know that the only changes in our models during adaptation with the induction gratings are restricted to the plasticity of lateral inhibitory connections and to homeostasis: in fact, learning in afferent projections and lateral excitatory connections was always disabled during our experiments, and every other variable in the model is constant.

What was left to determine was the individual contributions of the two mechanisms active during our experiments. It is interesting to note that using computer simulations provides a great advantage over physiological experiments as we can turn off either mechanism without affecting the rest of the system. This allows for predictions and a deeper investigation of the phenomenon, which could then hopefully be tested on the biological system.

Preliminary results comparing the strength of the effect computed for gratings oriented like the induction patterns hinted that homeostasis could account for the phenomenon, as it was found to induce strong and balanced effects. However, computing the orientation-dependent plot in each case (enabling only homeostatic adaptation or cortical plasticity) resulted in more interesting data. In fact, plots produced by inducing the McCollough Effect using only homeostatic mechanisms resulted in curves that were significantly narrower than those measured in human experiments, even though they were extremely accurate near the peaks (that is, at the orientations of the inducing patterns). On the contrary, simulation that only enabled cortical plasticity were found to be a perfect fit to the main characteristics of the psychological data, accounting for the width of the curves and the linear decay of the effect, which faded completely at 45^o off the orientation of the induction patterns.

There are two main reasons that could explain the good fit in the case of the homeostatic-based simulations and for the higher error around the peaks in the case of the plasticity-based ones. First, the low quality of the representation of the orientations in the model introduces errors in the plot, possibly activating similar sets of neurons regardless of the test patterns' orientations (e.g., near the orientation of the induction patterns). Second, our dichromatic model was characterized by a great number of double-duty cells (actually, almost every unit in the model was double-duty, see Fig. 3.2), in contrast to previous work [11]. A Primary Visual Cortex with many double-duty cells selective to many combinations of colors and orientations could give rise to a McCollough Effect, but its properties would likely be different from those observed in humans, and it would be unlikely to explain the long-term persistence of the effect anyway. Even in the case of a high number of double-duty cells, in fact, the width of the curve of the orientation-dependent plot would mostly depend on the width of the orientation tuning curve of the neurons and on the quality of the representation of the orientations.

It is worth noting that in the case of fewer double-duty neurons (as it is likely to be the case in primates [28]) we would expect the curves to be narrower when only homeostasis is enabled, and the overall effect to be described with worse accuracy than in the present simulations (i.e., the two plots, homeostasis-only and plasticity-only, Fig. 4.4, would be significantly different from each other, with the homeostatic-based one showing strong, narrow peaks at 0^o and 90^o and the plasticity-based one closely resembling the default simulation). Also, having a model with fewer double-duty units would probably imply a higher quality in the representation of orientations, as color and orientation would be learnt in a decorrelated fashion.

5.2 Future work

Future work will be required to better determine the basic requirements for reproducing the McCollough Effect, with special attention to the individual contributions of homeostatic mechanisms and cortical plasticity, which can be best studied using computer simulations.

Our models need to be improved in their representation of orientations and it would be useful to understand what does it drive the development and proportion of double-duty units in the models, as having the possibility to change their proportion would allow for further investigation on the role of homeostatic adaptation in the short-term dynamics of the McCollough Effect.

Also, physiological data suggests that the number of V1 neurons that are selective to color is rather small compared to the number of orientation-selective neurons [28], and this should be taken into account as a target for better, more biologically realistic, models.

It would also be interesting to produce a model (even "artificially", that is, not learnt) that had very narrow orientation tuning curves and a very accurate representation of orientations. In this case the homeostatic mechanisms alone would likely account for only a very little set of orientations close to the orientation of the induction gratings, and wouldn't probably produce a strong effect when other test gratings are used.

At last, more physiological experiments are required to better understand the representation of color in V1 and its relation to orientation, like determining the proportion and distribution of double-duty neurons and the quality of the representation of orientations.

5.2.1 Relevance to Neuroscience

Interestingly, our computational investigations can have broader applications than just the McCollough Effect. The technical possibility of separately handling the various contributions of adaptation and plasticity allows for a deep investigation of a range of cortical phenomena that apply to all visual experience, and not only to this specific phenomenon.

As we briefly mentioned, the selective disabling of either mechanism in the model while avoiding the disruption of any other cortical function is a key advantage of running computer simulations over physiological experiments. Both types of work are necessary, and data from studies on the direct biological system is critical to the development of accurate simulations. Having the possibility of running this type of simulations, however, can in turn aid the understanding of specific phenomena and can also suggest more

experiments to be performed.

The specific dynamics of the McCollough Effect that we simulated can be used not only to explain other types of after-effects, like it was done in previous studies, that however lacked homeostatic mechanisms [7, 11], but also a range of phenomena related to visual perception, like the learning of statistical properties of natural scenes [46] and the statistical decoding of feature properties from the Primary Visual Cortex.

Bibliography

[1] L. G. Allan and S. Siegel. McCollough Effects as conditioned responses: Reply to skowbo. 1986.

[2] B. Ans, C. Marendaz, J. Hérault, and B. Séré. McCollough Effect: a neural network model based on source separation. *Visual cognition*, 8(6):823–841, 2001.

[3] H. Barlow. A theory about the functional role and synaptic mechanism of visual after-effects. *Vision: Coding and efficiency*, 363375, 1990.

[4] H. B. Barlow. The knowledge used in vision and where it comes from. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 352(1358):1141–1147, 1997.

[5] J. A. Bednar. Topographica: Building and analyzing map-level simulations from Python, C/C++, MATLAB, NEST, or NEURON components. *Frontiers in Neuroinformatics*, 3:8, 2009.

[6] J. A. Bednar, J. B. De Paula, and R. Miikkulainen. Self-organization of color opponent receptive fields and laterally connected orientation maps. *Neurocomputing*, 65:69–76, 2005.

[7] J. A. Bednar and R. Miikkulainen. Tilt aftereffects in a self-organizing model of the primary visual cortex. *Neural computation*, 12(7):1721–1740, 2000.

[8] J. Broerse, T. Vladusich, and R. P. OShea. Colour at edges and colour spreading in McCollough Effects. *Vision Research*, 39(7):1305–1320, 1999.

[9] P. Buzás, P. Kóbor, Z. Petykó, I. Telkes, P. R. Martin, and L. Lénárd. Receptive field properties of color opponent neurons in the cat lateral geniculate nucleus. *The Journal of Neuroscience*, 33(4):1451–1461, 2013.

[10] P. Buzás, B. A. Szmajda, M. Hashemi-Nezhad, B. Dreher, and P. R. Martin. Color signals in the primary visual cortex of marmosets. *Journal of vision*, 8(10):7, 2008.

[11] J. B. Ciroux. Simulating the McCollough Effect in a self-organizing model of the primary visual cortex. Master's thesis, The University of Edinburgh, 2006.

[12] B. R. Conway. Colour vision: A clue to hue in V2. *Current Biology*, 13(8):R308 – R310, 2003.

[13] F. Crick and C. Koch. Are we aware of neural activity in primary visual cortex? *Nature*, 375(6527):121–123, 1995.

[14] P. C. Dodwell and G. K. Humphrey. A functional theory of the McCollough Effect. *Psychological review*, 97(1):78, 1990.

[15] F. H. Durgin and D. R. Proffit. Visual learning in the perception of texture: simple and contingent aftereffects of texture density. *Spatial vision*, 9(4):423–474, 1996.

[16] S. R. Ellis. Orientation selectivity of the McCollough Effect: Analysis by equivalent contrast transformation. *Perception & Psychophysics*, 22(6):539–544, 1977.

[17] L. Fidell. Orientation specificity in chromatic adaptation of human "edge-detectors". *Perception & Psychophysics*, 8(4):235–237, 1970.

[18] J. D. Forte and C. W. Clifford. Inter-ocular transfer of the tilt illusion shows that monocular orientation mechanisms are colour selective. *Vision research*, 45(20):2715–2721, 2005.

[19] W. S. Geisler and D. G. Albrecht. Visual cortex neurons in monkeys and cats: detection, discrimination, and identification. *Visual neuroscience*, 14(05):897–919, 1997.

[20] S. Grossberg, S. Hwang, and E. Mingolla. Thalamocortical dynamics of the McCollough Effect: boundary-surface alignment through perceptual learning. *Vision research*, 42(10):1259–1286, 2002.

[21] R. Held and S. R. Shattuck. Color-and edge-sensitive channels in the human visual system: tuning for orientation. *Science*, 174(4006):314–316, 1971.

[22] G. K. Humphrey and M. A. Goodale. Probing unconscious visual processing with the McCollough Effect. *Consciousness and cognition*, 7(3):494–519, 1998.

[23] L. M. Hurvich and D. Jameson. An opponent-process theory of color vision. *Psychological review*, 64(6p1):384, 1957.

[24] K. Jameson, S. Highnote, and L. Wasserman. Richer color experience in observers with multiple photopigment opsin genes. *Psychonomic Bulletin & Review*, 8(2):244–261, 2001.

[25] E. N. Johnson, M. J. Hawken, and R. Shapley. The orientation selectivity of color-responsive neurons in macaque v1. *The Journal of Neuroscience*, 28(32):8096–8106, 2008.

[26] E. N. Johnson, S. D. Van Hooser, and D. Fitzpatrick. The representation of s-cone signals in primary visual cortex. *The Journal of Neuroscience*, 30(31):10337–10350, 2010.

[27] P. D. Jones and D. H. Holding. Extremely long-term persistence of the McCollough Effect. *Journal of Experimental Psychology: Human Perception and Performance*, 1(4):323, 1975.

[28] C. E. Landisman and D. Y. Ts'o. Color processing in macaque striate cortex: relationships to ocular dominance, cytochrome oxidase, and orientation. *Journal of Neurophysiology*, 87(6):3126–3137, 2002.

[29] H. D. Lu and A. W. Roe. Functional organization of color domains in V1 and V2 of macaque monkey revealed by optical imaging. *Cerebral Cortex*, 18(3):516–533, 2008.

[30] D. MacKay and V. MacKay. Retention of the McCollough Effect in darkness: Storage or enhanced read-out? *Vision research*, 17(2):313–315, 1977.

[31] C. McCollough. Color adaptation of edge-detectors in the human visual system. *Science*, 149(3688):1115–1116, 1965.

[32] C. R. Michael. Color vision mechanisms in monkey striate cortex: simple cells with dual opponent-color receptive fields. *Journal of Neurophysiology*, 1978.

[33] R. Miikkulainen, J. A. Bednar, Y. Choe, and J. Sirosh. *Computational maps in the visual cortex*, volume 14. Springer, 2005.

[34] T. Morita, T. Kochiyama, T. Okada, Y. Yonekura, M. Matsumura, and N. Sadato. The neural substrates of conscious color perception demonstrated using fMRI. *Neuroimage*, 21(4):1665–1673, 2004.

[35] V. M. Parraga C. A., Baldrich R. Accurate mapping of natural scenes radiance to cone activation space: A new image dataset. *CGIV 2010/MCS'10*, 2010.

[36] J. B. D. Paula. *Modeling the self-organization of color selectivity in the visual cortex*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX, 2007.

[37] J. Pokorny and V. C. Smith. Psychophysical signatures associated with magnocellular and parvocellular pathway contrast gain. *JOSA A*, 14(9):2477–2486, 1997.

[38] L. A. Riggs, K. D. White, and P. D. Eimas. Establishment and decay of orientation-contingent aftereffects of color. *Perception & Psychophysics*, 16(3):535–542, 1974.

[39] J.-L. R. Stevens, J. S. Law, J. Antolík, and J. A. Bednar. Mechanisms for stable, robust, and adaptive development of orientation maps in the primary visual cortex. *The Journal of Neuroscience*, 33(40):15747–15766, 2013.

[40] C. Stromeyer. Further studies of the McCollough Effect. *Perception & Psychophysics*, 6(2):105–110, 1969.

[41] P. Thompson and D. Burr. Visual aftereffects. *Current Biology*, 19(1):R11–R14, 2009.

[42] P. Thompson and G. Latchford. Colour-contingent after-effects are really wavelength-contingent. 1986.

[43] E. Vul, E. Krizay, and D. I. MacLeod. The McCollough Effect reflects permanent and transient adaptation in early visual cortex. *Journal of Vision*, 8(12):4, 2008.

[44] E. Vul and D. I. MacLeod. Contingent aftereffects distinguish conscious and preconscious color processing. *Nature neuroscience*, 9(7):873–874, 2006.

[45] M. A. Webster. Adaptation and visual coding. *Journal of vision*, 11(5):3, 2011.

[46] M. A. Webster and J. Mollon. Adaptation and the color statistics of natural images. *Vision research*, 37(23):3283–3298, 1997.

[47] W. Webster, R. Day, and K. Willenberg. Orientation-contingent color aftereffects are determined by real color, not induced color. *Perception & psychophysics*, 44(1):43–49, 1988.

[48] Y. Xiao, A. Casti, J. Xiao, and E. Kaplan. Hue maps in primate striate cortex. *Neuroimage*, 35(2):771–786, 2007.

[49] J. A. Yamashita, J. L. Hardy, K. K. De Valois, and M. A. Webster. Stimulus selectivity of figural aftereffects for faces. *Journal of Experimental Psychology: Human Perception and Performance*, 31(3):420, 2005.