UNIVERSITÀ DEGLI STUDI DI TRIESTE

Sede Amministrativa del Dottorato di Ricerca

XXI Ciclo del
Dottorato di Ricerca in
Ingegneria dell'Informazione

HIGH DYNAMIC RANGE IMAGES: PROCESSING, DISPLAY AND PERCEPTUAL QUALITY ASSESSMENT

(Settore scientifico-disciplinare ING-INF/01)

Dottorando:
Dott. Ing. Gabriele Guarnieri

Coordinatore del Collegio Docenti:
Chiar.mo Prof. Roberto Vescovo
Università di Trieste

Tutore e Relatore:
Chiar.mo Prof. Giovanni Ramponi
Università di Trieste

Correlatore:
Dott. Ing. Luigi Albani
FIMI-Philips, Saronno

ANNO ACCADEMICO 2007/2008

I believe in things that are developed through hard work. I always like people who have developed long and hard, especially through introspection and a lot of dedication. I think what they arrive at is usually a much deeper and more beautiful thing than the person who seems to have that ability and fluidity from the beginning.

Bill Evans

A technical note

The PDF version of this thesis was compiled with a particular care towards the accurate reproduction of the gray levels and the colors in the images. Where possible, the images were generated using appropriate color management algorithms and the ICC profiles of the corresponding color space were embedded into the PDF document.

For a best rendition, this document should be viewed on a calibrated display with a PDF reader capable of loading the ICC profile of the monitor and performing a color management.

Introduction

The intensity of natural light can span over 10 orders of magnitude from starlight to direct sunlight. Even in a single scene, the luminance of the bright areas can be thousands or millions of times greater than the luminance in the dark areas; the ratio between the maximum and the minimum luminance values is commonly known as *dynamic range* or *contrast*. The human visual system is able to operate in an extremely wide range of luminance conditions without saturation and at the same time it can perceive fine details which involve small luminance differences. Our eyes achieve this ability by modulating their response as a function of the local mean luminance with a process known as *local adaptation*. In particular, the visual sensation is not linked to the absolute luminance, but rather to its spatial and temporal variation. One consequence of the local adaptation capability of the eye is that the objects in a scene maintain their appearance even if the light source illuminating the scene changes significantly. On the other hand, the technologies used for the acquisition and reproduction of digital images are able to handle correctly a significantly smaller luminance range of 2 to 3 orders of magnitude at most. Therefore, a *high dynamic range* (HDR) image poses several challenges and requires the use of appropriate techniques. These elementary observations define the context in which the entire research work described in this Thesis has been performed. As indicated below, different fields have been considered; they range from the acquisition of HDR images to their display, from visual quality evaluation to medical applications, and include some developments on a recently proposed class of display equipment.

An HDR image can be captured by taking multiple photographs with different exposure times or by using high dynamic range sensors; moreover, synthetic HDR images can be generated with computer graphics by means of physically-based algorithms which often involve advanced lighting simulations. An HDR image, although acquired correctly, can not be displayed on a conventional monitor. The white level of most devices is limited to a few hundred cd/m^2 by technological constraints, primarily linked to the power consumption and heat dissipation; the black level also has a non negligible luminance, in particular for devices based on the liquid crystal technology. However, thanks to the aforementioned properties of the human visual system, an exact reproduction of the luminance in the original scene is not strictly necessary in order to produce a similar sensation in the observer. For this purpose, *dynamic range reduction* algorithms have been developed which attenuate the large luminance variations in an image while preserving as far as possible the fine details.

The most simple dynamic range reduction algorithms map each pixel individually with the same nonlinear function commonly known as *tone mapping curve*. One operator we propose, based on a modified logarithmic function, has a low computational cost and contains one single user-adjustable parameter. However, the methods belonging to this category can reduce the visibility of the details in some portions of the image. More advanced methods also take into account

the pixel neighborhood. This approach can achieve a better preservation of the details, but the loss of one-to-one mapping from input luminances to display values can lead to the formation of gradient reversal effects, which typically appear as halos around the object boundaries. Different solutions to this problem have been attempted. One method we introduce is able to avoid the formation of halos and intrinsically prevents any clipping of the output display values. The method is formulated as a constrained optimization problem, which is solved efficiently by means of appropriate numerical methods.

In specific applications, such as the medical one, the use of dynamic range reduction algorithms is discouraged because any artifacts introduced by the processing can lead to an incorrect diagnosis. In particular, a one-to-one mapping from the physical data (for instance, a tissue density in radiographic techniques) to the display value is often an essential requirement. For this purpose, high dynamic range displays, capable of reproducing images with a wide luminance range and possibly a higher bit depth, are under active development. Dual layer LCD displays, for instance, use two liquid crystal panels stacked one on top of the other over an enhanced backlight unit in order to achieve a dynamic range of $4 \div 5$ orders of magnitude. The grayscale reproduction accuracy is also increased, although a "bit depth" can not be defined unambiguously because the luminance levels obtained by the combination of the two panels are partially overlapped and unevenly spaced. A dual layer LCD display, however, requires the use of complex *splitting algorithms* in order to generate the two images which drive the two liquid crystal panels. A splitting algorithm should compensate multiple sources of error, including the parallax introduced by the viewing angle, the gray-level clipping introduced by the limited dynamic range of the panels, the visibility of the reconstruction error, and glare effects introduced by an unwanted light scattering between the two panels. For these reasons, complex constrained optimization techniques are necessary. We propose an objective function which incorporates all the desired constraints and requirements and can be minimized efficiently by means of appropriate techniques based on multigrid methods.

The quality assessment of high dynamic range images requires the development of appropriate techniques. By their own nature, dynamic range reduction algorithms change the luminance values of an image significantly and make most image fidelity metrics inapplicable. Some particular aspects of the methods can be quantified by means of appropriate operators; for instance, we introduce an expression which describes the detail attenuation introduced by a tone mapping curve. In general, a subjective quality assessment is preferably performed by means of appropriate psychophysical experiments. We conducted a set of experiments, targeted specifically at measuring the level of agreement between different users when adjusting the parameter of the modified logarithmic mapping method we propose. The experimental results show a strong correlation between the user-adjusted parameter and the image statistics, and suggest a simple technique for the automatic adjustment of this parameter. On the other hand, the quality assessment in the medical field is preferably performed by means of objective methods. In particular, task-based quality measures evaluate by means of appropriate observer studies the clinical validity of the image used to perform a specific diagnostic task. We conducted a set of observer studies following this approach, targeted specifically at measuring the clinical benefit introduced by a high dynamic range display based on the dual layer LCD technology over a conventional display with a low dynamic range and 8-bit quantization. Observer studies are often time consuming and difficult to organize; in order to increase the number of tests, the human observers can be partially replaced by appropriate software applications, known as *model observers* or *computational observers*, which simulate the diagnostic task by means of statistical classification techniques.

This thesis is structured as follows. Chapter 1 contains a brief background of concepts related to the physiology of human vision and to the electronic reproduction of images. The description we make is by no means complete and is only intended to introduce some concepts which will be extensively used in the following. Chapter 2 describes the technique of high dynamic range image acquisition by means of multiple exposures. In Chapter 3 we introduce the dynamic range reduction algorithms, providing an overview of the state of the art and proposing some improvements and novel techniques. In Chapter 4 we address the topic of quality assessment in dynamic range reduction algorithms; in particular, we introduce an operator which describes the detail attenuation introduced by tone mapping curves and describe a set of psychophysical experiments we conducted for the adjustment of the parameter in the modified logarithmic mapping method we propose. In Chapter 5 we move to the topic of medical images and describe the techniques used to map the density data of radiographic images to display luminances. We point out some limitations of the current technical recommendation and propose an improvement. In Chapter 6 we describe in detail the dual layer LCD prototype and propose different splitting algorithms for the generation of the two images which drive the two liquid crystal panels. In Chapter 7 we propose one possible technique for the estimation of the equivalent bit depth of a dual layer LCD display, based on a statistical analysis of the quantization noise. Finally, in Chapter 8 we address the topic of objective quality assessment in medical images and describe a set of observer studies we conducted in order to quantify the clinical benefit introduced by a high dynamic range display.

No general conclusions are offered; the breadth of the subjects has suggested to draw more focused comments at the end of the individual chapters.

Introduzione

L'intensità della luce naturale può coprire oltre 10 ordini di grandezza dalla luce stellare alla luce solare diretta. Anche in una singola scena, la luminanza nelle zone chiare può essere migliaia o milioni di volte superiore alla luminanza nelle zone scure; il rapporto tra la luminanza massima e minima viene comunemente detto *dinamica* o *contrasto*. Il sistema visivo umano è in grado di funzionare senza saturazione in condizioni di luce estremamente ampie e allo stesso tempo di percepire dettagli fini che comportano piccole differenze di luminosità. I nostri occhi riescono a raggiungere questo scopo modulando la loro risposta in funzione della luminanza media locale con un processo noto come *adattamento locale*. In particolare, la sensazione visiva non è determinata dalla luminanza assoluta, bensì alla sua variazione spaziale e temporale. Una conseguenza della capacitò di adattamento locale dell'occhio è che gli oggetti in una scena mantengono quasi inalterato il loro aspetto anche se la sorgente di luce che illumina la scena cambia significativamente. Tuttavia, gli strumenti tecnologici utilizzati per l'acquisizione e la riproduzione di immagini digitali sono in grado di gestire correttamente un intervallo di luminanze molto più ristretto, di $2 \div 3$ ordini di grandezza al massimo. Pertanto, un'immagine ad *alta dinamica* pone numerosi problemi e richiede l'uso di tecniche apposite. Queste osservazioni definiscono a grandi linee l'ambito nel quale si è svolto l'intero lavoro di ricerca descritto in questa Tesi. Come indicato di seguito, si sono considerati diversi campi; essi spaziano dall'acquisizione di immagini ad alta dinamica alla loro visualizzazione, dalla valutazione della qualità visiva alle applicazioni medicali, e comprendono alcuni studi su una tipologia di schermi proposta recentemente.

Un'immagine ad alta dinamica può essere acquisita scattando più fotografie con diversi tempi di esposizione oppure usando appositi sensori ad alta dinamica; inoltre, si possono generare immagini sintetiche con tecniche di grafica computerizzata che utilizzano complesse simulazioni delle proprietà fisiche della luce. Un'immagine ad alta dinamica, anche se acquisita correttamente, non può essere visualizzata su un monitor convenzionale. Il livello del bianco nella maggior parte dei dispositivi è limitato a poche centinaia di cd/m^2 da vincoli tecnologici legati soprattutto al consumo di potenza e alla dissipazione del calore; il livello del nero ha anche una luminanza non trascurabile , in particolare nei dispositivi basati sulla tecnologia a cristalli liquidi. Tuttavia, grazie alle proprietà dell'occhio umano citate in precedenza, non è necessario riprodurre esattamente la luminanza della scena reale per indurre una simile sensazione nell'osservatore. A questo scopo sono stati sviluppati algoritmi per la riduzione della dinamica, che attenuano le grandi variazioni di luce in un'immagine preservando nei limiti del possibile i dettagli più fini.

Gli algoritmi più semplici per la riduzione della dinamica mappano ciascun pixel individualmente con un'unica funzione non lineare comunemente nota come *curva di mappatura*. Un operatore che proponiamo, basato su una funzione logaritmica modificata, presenta un

basso costo computazionale e contiene un singolo parametro da regolare. Tuttavia, i metodi appartenenti a questa categoria possono ridurre la visibilità dei dettagli in alcune porzioni dell'immagine. I metodi più avanzati tengono anche conto dei pixel circostanti. Questa tecnica permette di conservare meglio i dettagli, ma la perdita di corrispondenza biunivoca tra luminanze in ingresso e valori visualizzati può causare effetti di inversione del gradiente, che tipicamente si manifestano sotto forma di aloni intorno ai bordi degli oggetti. Si sono tentate diverse soluzioni a questo problema. Un metodo che proponiamo è in grado di evitare la formazione di aloni e previene intrinsecamente ogni tipo di saturazione nei valori visualizzati. Il metodo è formulato come un problema di ottimizzazione vincolata, che viene risolto in modo efficiente mediante apposite tecniche di calcolo numerico.

In particolari applicazioni, come quella medicale, si sconsiglia l'uso di algoritmi per la riduzione della dinamica perché qualunque artefatto introdotto dall'elaborazione potrebbe portare ad una diagnosi errata. In particolare, la corrispondenza biunivoca dai dati fisici (ad esempio la densità dei tessuti nel caso di tecniche radiografiche) e valori visualizzati è spesso un requisito essenziale. Per questo scopo, si stanno attivamente sviluppando monitor ad alta dinamica in grado di riprodurre immagini con un ampio intervallo di valori di luminosità e possibilmente una profondità di bit maggiore. Ad esempio, i monitor a cristalli liquidi a doppio strato usano due pannelli a cristalli liquidi, sovrapposti uno sull'altro di fronte ad un'unità di retroilluminazione potenziata, per ottenere una dinamica di $4 \div 5$ ordini di grandezza. L'accuratezza nella riproduzione dei livelli è anche migliorata, anche se non si può indicare in modo univoco la profondità di bit perché i livelli ottenuti dalla combinazione dei due pannelli sono parzialmente sovrapposti e distribuiti in modo irregolare. Un monitor a doppio pannello, tuttavia, richiede l'uso di complessi *algoritmi di scomposizione* per generare le due immagini che pilotano i due pannelli a cristalli liquidi. Un algoritmo di scomposizione dovrebbe compensare molteplici fonti di errore, tra cui il parallasse introdotto dall'angolo di visione, la saturazione causata dalla dinamica limitata dei pannelli, la visibilità dell'errore di ricostruzione e gli artefatti introdotti da diffusioni indesiderate della luce tra i due pannelli. Per questi motivi, è necessario utilizzare complesse tecniche di ottimizzazione vincolata. Proponiamo una funzione obiettivo che incorpora tutti i vincoli e i requisiti di progetto e che può essere minimizzata in modo efficiente con apposite tecniche basate sui metodi "multigrid".

La valutazione della qualità di immagini ad alta dinamica richiede lo sviluppo di tecniche apposite. Per la loro stessa natura, gli algoritmi di riduzione della dinamica alterano significativamente i valori di luminosità in un'immagine rendendo inapplicabile la maggior parte delle misure di fedeltà. Alcuni particolari aspetti dei metodi possono essere quantificati mediante operatori appositi; ad esempio proponiamo un'espressione che descrive l'attenuazione dei dettagli introdotta da una curva di mappatura. In generale, è preferibile effettuare una valutazione soggettiva della qualità mediante appositi esperimenti psicofisici. Abbiamo condotto una serie di esperimenti, con l'obiettivo specifico di valutare il livello di accordo tra diversi utenti nella regolazione del parametro della mappatura logaritmica modificata da noi proposta. I risultati sperimentali mostrano una forte correlazione tra il valore del parametro regolato dall'utente e le proprietà statistiche dell'immagine, e suggeriscono una semplice tecnica per la regolazione automatica del parametro in questione. La valutazione della qualità nel campo medico, invece, si effettua preferibilmente mediante metodi oggettivi. In particolare, le misure di qualità "task-based" valutano in appositi studi il valore clinico delle immagini utilizzate per uno specifico compito diagnostico. Abbiamo realizzato una serie di prove di questo tipo, con l'obiettivo specifico di misurare il beneficio clinico portato da un monitor ad alta dinamica, basato sulla tecnologia a cristalli liquidi a doppio strato, rispetto ad un monitor

convenzionale a bassa dinamica e con una quantizzazione a 8 bit. Le prove con osservatori comportano spesso tempi lunghi e sono difficili da organizzare; per aumentare il numero di prove, si può sostituire in parte gli osservatori umani con appositi programmi software noti come "model observers" o "computational observers", che simulano l'operazione di diagnosi mediante tecniche di classificazione statistica.

Questa Tesi è strutturata nel modo seguente. Il Capitolo 1 contiene un breve richiamo di concetti riguardanti la fisiologia della visione e la riproduzione elettronica delle immagini. La trattazione che facciamo non ha pretese di completezza e ha il solo scopo di introdurre alcuni concetti che verranno usati frequentemente in seguito. Il Capitolo 2 descrive il metodo di acquisizione di immagini ad alta dinamica mediante esposizioni multiple. Nel Capitolo 3 introduciamo gli algoritmi per la riduzione della dinamica, fornendo una descrizione dello stato dell'arte e proponendo alcuni miglioramenti e alcune tecniche innovative. Nel Capitolo 4 trattiamo l'argomento della valutazione della qualità dei metodi per la riduzione della dinamica; in particolare, proponiamo un operatore che descrive l'attenuazione dei dettagli prodotta dalle curve di mappatura e descriviamo una serie di esperimenti psicofisici che abbiamo realizzato per la regolazione del parametro della mappatura logaritmica modificata da noi proposta. Nel Capitolo 5 passiamo al tema delle immagini medicali e descriviamo le tecniche usate per mappare la densità dei tessuti nelle immagini radiografiche in luminanza visualizzata. Indichiamo alcune limitazioni della normativa attualmente in vigore e proponiamo un miglioramento. Nel Capitolo 6 descriviamo in dettaglio il prototipo di monitor a cristalli liquidi a doppio strato e proponiamo diversi algoritmi di scomposizione per la generazione delle immagini che pilotano i due pannelli a cristalli liquidi. Nel Capitolo 7 proponiamo una possibile tecnica per la stima della profondità di bit equivalente di un monitor a doppio pannello, basata su un'analisi statistica del rumore di quantizzazione. Infine, nel Capitolo 8 trattiamo l'argomento della valutazione oggettiva della qualità nelle immagini mediche e descriviamo una serie di prove che abbiamo realizzato per quantificare il beneficio clinico introdotto da un monitor ad alta dinamica.

Non verranno tratte delle conclusioni generali; vista l'ampiezza degli argomenti trattati, abbiamo preferito riportare dei commenti specifici alla fine di ciascun capitolo.

Contents

Chapter 1

Some aspects of human vision and artificial image reproduction

1.1 Summary

This chapter contains a brief background to concepts related both to the physiology of human vision, including color and luminance perception, and to the electronic reproduction of images, including their acquisition by means of digital cameras and their visualization on display devices. We also describe the most widespread standard recommendations used in the measurement of luminance and color. We chose to describe the three topics of vision, image acquisition and image reproduction in the same chapter because their structure is closely related; in particular, the technique used by the display devices to reproduce color is made possible by the particular structure of the light-sensitive cells in the human eye.

The aforementioned subjects are complex and, in some aspects, not yet fully understood. Their description in this chapter is by no means complete. Indeed, our goal is neither to provide a comprehensive treatise on the subject, nor to introduce novel contributions, but only to introduce some concepts which will be extensively used in this thesis. We also believe that it may be beneficial to provide a reference for some concepts which are often misunderstood and erroneously described.

1.2 Physiology of color vision

The retina in the human eye contains two categories of photoreceptor cells [80]. The *rod* cells are approximately 120 millions, mainly located at the sides of the retina. Rod cells are sensitive to low light intensities, being capable to detect in some cases even one single photon, and are therefore involved in the night or *scotopic* vision. The *cone* cells are approximately 6 millions, mainly located in the central portion of the retina known as *fovea*. Cone cells have a lower sensitivity and are therefore responsible for daylight or *photopic* vision. The human eye contains three types of cone cells, sensitive to long (L), medium (M) and short (S) wavelengths which roughly correspond to the red, green and blue colors of visible light. This allows the eye to discriminate, to a certain extent, the wavelength of the incident light, thus enabling color vision.

The plot in Figure 1.1 shows the results of an actual measurement of the spectral sensitivity of the cones and rods in the human retina [38].

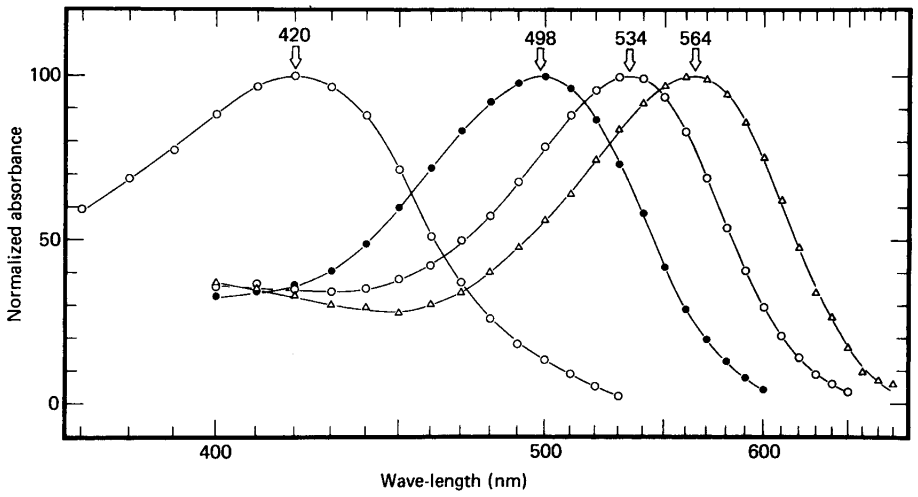

Figure 1.1: Normalized spectral sensitivity (from left to right) of the S cones, rods, M cones and L cones in the human retina. Reproduced from [38].

The light which reaches the retina induces some chemical reactions on the proteins contained in the cell pigments, and these reactions generate an electric potential across the cell membrane. Let $I(\lambda)$ indicate the power spectrum of the incident light and $\rho_L(\lambda)$, $\rho_M(\lambda)$ and $\rho_S(\lambda)$ the spectral sensitivities of the three types of cone cells. The response of the photosensitive cells is described by the three expressions

$$L \triangleq \int \rho_L(\lambda)\, I(\lambda)\, d\lambda \qquad M \triangleq \int \rho_M(\lambda)\, I(\lambda)\, d\lambda \qquad S \triangleq \int \rho_S(\lambda)\, I(\lambda)\, d\lambda. \tag{1.1}$$

The physical stimulus involved in *trichromatic* color vision consists therefore of three scalar values and carries little information about the spectrum of the incident light. In other words, the mapping is not injective and different spectra can induce the same visual stimulus. This property is known as *metamerism* and plays a fundamental role in the artificial reproduction of color. In order to reproduce – for instance on a cathode ray tube (CRT) or liquid crystal display (LCD) – the color sensation produced by the observation of a real-life scene, it is not necessary to generate a light with identical spectral characteristics; it is sufficient that the LMS responses (1.1) of the three types of cone cells be the same in both viewing conditions. One simple technique, which allows to reproduce a wide gamut of colors, consists in mixing three predetermined light sources commonly known as *primaries*, adjusting their intensities in order to achieve the match. The spectrum of the three primaries is arbitrary, with the only constraint of linear independence; however, red, green and blue (RGB) primaries are most commonly used for a reason which will be clear in the following.

Let $I_R(\lambda)$, $I_G(\lambda)$ and $I_B(\lambda)$ indicate the normalized power spectra of the red, green and blue primaries and R, G and B their intensities. A linear combination of the primaries reproduces the original spectrum $I(\lambda)$ if the following condition is satisfied:

$$R\, I_R(\lambda) + G\, I_G(\lambda) + B\, I_B(\lambda) \overset{\text{metamer}}{=\!=\!=} I(\lambda) \tag{1.2}$$

By multiplying the left and right hand sides of Equation (1.2) with the spectral sensitivities of the cone cells and integrating over the wavelength λ, the following linear system is obtained:

$$\begin{cases} R\int \rho_L(\lambda)\, I_R(\lambda)\, d\lambda + G\int \rho_L(\lambda)\, I_G(\lambda)\, d\lambda + B\int \rho_L(\lambda)\, I_B(\lambda)\, d\lambda = \int \rho_L(\lambda)\, I(\lambda)\, d\lambda \\ R\int \rho_M(\lambda)\, I_R(\lambda)\, d\lambda + G\int \rho_M(\lambda)\, I_G(\lambda)\, d\lambda + B\int \rho_M(\lambda)\, I_B(\lambda)\, d\lambda = \int \rho_M(\lambda)\, I(\lambda)\, d\lambda \\ R\int \rho_S(\lambda)\, I_R(\lambda)\, d\lambda + G\int \rho_S(\lambda)\, I_G(\lambda)\, d\lambda + B\int \rho_S(\lambda)\, I_B(\lambda)\, d\lambda = \int \rho_S(\lambda)\, I(\lambda)\, d\lambda \end{cases} \quad (1.3)$$

If the three primaries are linearly independent, the system (1.3) has a unique solution (R, G, B) which provides the intensity of the three primaries which reproduces the original color. These values are often improperly used as a definition of the color itself, although they do not describe the physical nature of color but only one particular technique used for its reproduction. In particular, the RGB values do not correspond to the LMS cone cell responses, but are only related to them by means of a linear transformation which depends on the particular choice of the primaries.

The use of three primaries is motivated by the fact that the human eye contains three types of cone cells: in particular, this correspondence guarantees that the system (1.3) has the same number of equations and unknowns. However, this choice does not allow to reproduce any visible color. By their physical nature, the RGB intensities can only take positive values; therefore, a color can be reproduced only if the linear system (1.3) has a positive solution. The set of colors which can be reproduced using a given set of primaries is commonly known as *gamut*. The use of red, green and blue primaries – which roughly correspond to wavelengths at the two ends and at the middle of the visible range – is the most widespread, since it provides a gamut which contains most of the colors in natural scenes. In professional applications, *extended color gamut* displays with four or more primaries are sometimes used.

1.3 Chromatic adaptation and color constancy

The electric signal generated by the photosensitive cells in the human retina as a response to a light stimulus is, to a good approximation, a linear function of the stimulus intensity described by the expressions (1.1). This signal is then transmitted through different layers of cells known as the *horizontal*, *bipolar* and *amacrine* cells, and finally to the *ganglion* cells which form the optical nerve [39]. These layers of cells introduce a complex, and not yet fully understood, processing of the electric signal which includes spatial interactions among adjacent cells, nonlinear mappings and feedback mechanisms. As a result of these mechanisms, the signal which reaches the brain through the optical nerve is a nonlinear function of the light intensity which does not simply depend on the characteristics of the fixated point but is strongly influenced by the surround.

A notable set of experiments in this field was conducted in the 1960s by Edwin Land, founder of Polaroid corporation, and later by John McCann [44, 43]. In the experiments, two identical test targets were built by arranging about 100 sheets of paper of different colors. The papers had different shapes and sizes in order to inhibit the formation of afterimages on the observer's retina. The test targets were known under the nickname of "Mondrians" due to their resemblance with a work by the Dutch painter in the Tate Gallery of London. In the experiments, each of the two Mondrians was illuminated by three adjustable monochromatic projectors with long, medium and short wavelengths. Clearly, if the two Mondrians were identically illuminated, their colors also appeared identical to the observer. Land then adjusted

the projectors for the two Mondrians in such a way that two selected patches – for instance, a green paper on the first Mondrian and a red paper on the second one – had the same radiance for each spectral band. In this way, the two patches produced the same physical stimulus to the observer's eye. However, the observer still perceived the first patch as green and the second patch as red. By repeating this procedure with different pairs of patches, Land proved that the same physical stimulus can correspond to virtually any perceived color.

The described mechanism is known under different names, the most common ones being *chromatic adaptation* and *color constancy*. The color sensation perceived by the observer is not simply related to the pointwise characteristics of the fixated image, but rather to its spatial and, to some extent, temporal variation. One major consequence of this property is that the objects maintain their perceived color under different light sources, although the spectral characteristics of the light source can change within ample margins. Apparently, the interactions between the retinal cells have the effect of adjusting the gain for the three LMS bands individually and locally in order to attenuate or discard a color cast introduced by the spectral characteristics of the light source. One popular model, proposed by Land and McCann following the aforementioned Mondrian experiments, is based on the observation that the light intensity which reaches the eye is given by the product of *illumination*, which depends on the light source illuminating the scene, and *reflectance*, which instead depends on the surfaces of the objects. As a consequence of the chromatic adaptation mechanism, the perceived visual sensation is determined almost entirely by the reflectance, which is an intrinsic property of the object surfaces, and is influenced only minimally by the illumination, which can vary strongly under different light sources. Apparently, this processing takes place both in the retinal cells and in the visual cortex inside the brain; for this reason, the model was called *Retinex* by contraction of the two words *retina* and *cortex*.

The validity of the Retinex model is disputed. In particular, the original Mondrian experiments used a uniform diffuse illumination which is not a realistic model for natural scenes; the illumination in natural scenes can span a significantly larger dynamic range and the effect of the illumination becomes visible. In a recent experiment [56], a 3D Mondrian was built by assembling solids of different shapes and colors; the scene was then illuminated with either a diffuse light, resulting in a low dynamic range scene, or a directional spotlight, resulting in a high dynamic range scene. The experiment showed that, when the dynamic range of the scene is increased, the color constancy mechanism fails and the perceived color of the objects changes according to the illumination. Nevertheless, the Retinex model is useful and can be used profitably in computational algorithms which will be described in the following chapters.

1.4 Photometry

The intensity of light can be easily measured by its power. However, the human eye is sensitive only to a limited range of wavelengths, approximately ranging from 350 to 700 nm, and also in this interval the sensitivity strongly depends on the wavelength. Therefore, lights with the same physical power but different wavelengths can have a significantly different perceptual brightness for a human observer or be invisible at all. For this reason, in the measurement of visible light it is advisable to use photometric units which take into account the characteristics of human vision.

In 1924, the *Commission Internationale de l'Eclairage* (CIE) defined a curve which describes the

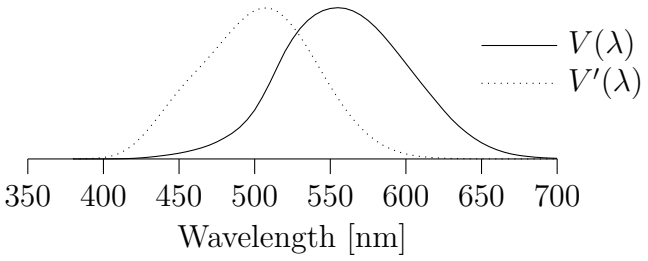

Figure 1.2: Photopic efficiency $V(\lambda)$ and scotopic efficiency $V'(\lambda)$ of the human eye.

sensitivity of the human eye in daylight as a function of the wavelength λ of the incident light. The curve is commonly known as *photopic efficiency* and is indicated with the symbol $V(\lambda)$. A different curve, introduced in 1951, describes the eye sensitivity at low levels of illumination when the rod cells are mainly used. This curve is known as *scotopic efficiency* and is indicated with the symbol $V'(\lambda)$. A plot of the two curves is shown in Figure 1.2. The photopic efficiency curve indicates that, for the same physical power, a green light with a wavelength of 555 nm appears the brightest; a red or blue light appears dimmer, and infrared or ultraviolet light outside the visible range is not visible. The light intensity corresponding to a power spectrum $I(\lambda)$ is then defined by the expression

$$L \triangleq K_m \int V(\lambda)\, I(\lambda)\, d\lambda \tag{1.4}$$

The symbol used to indicate the light intensity is not universally agreed upon; in this Thesis, we shall use the letter L. The units of measure depend on the measurement conditions. If we consider the total luminous flux emitted by a light source, the power spectrum $I(\lambda)$ is measured in W/m (watts per unit wavelength) and the light intensity L is measured in *lumen* (Lm). If instead we consider the light intensity falling on a surface, the power spectrum $I(\lambda)$ is measured in $\mathrm{W/m^3}$ (watts per unit area and unit wavelength) and the light intensity L is measured in *lux* (Lx). The scaling factor K_m has a value of 683 Lm/w and is used to convert between the physical and the photometric units. Its value was chosen for historical reasons. The base unit of luminous intensity in the International System of Units is the *candela* (cd), with a definition based on the luminance of a blackbody radiator at the melting temperature of platinum. The scaling factor K_m was chosen in such a way that one candela corresponds to one lumen per steradian. The light intensity emitted by a surface such as a display device is generally measured in candelas per square meter or *nits*. Although used in different contexts, lux and $\mathrm{cd/m^2}$ have the same physical dimension and the conversion factor is $1\ \mathrm{cd/m^2} = 4\pi$ lx.

1.5 Color matching experiments and the CIE-XYZ color space

Since the color sensation produced by the human eye is given by the responses of the three groups of cone cells in the retina, sensitive to the long (L), medium (M) and short (S) wavelengths, the most natural method of measuring the color correspondent to a light stimulus is to compute its LMS coordinates using the expression (1.1). However, such a definition poses some practical issues: in particular, the cone responses can vary among different observers,

and their measurement requires technologies which have become available only in recent times [38]. In 1931, the *Commission Internationale de l'Eclairage* (CIE) proposed a standard recommendation based on an indirect, approximate measurement of the eye sensitivity. The measurement is based on *color matching* experiments, in which an observer tries to reproduce, or *match*, the color of a test light source by adjusting the intensity of three primary light sources. This kind of experiment was proposed in 1853 by Hermann Grassmann, but the first rigorous and complete sessions were conducted in the 1920s by W. David Wright and John Guild [79, 31]. As mentioned in Section 1.2, the color of the three primaries is arbitrary with the only constraint of linear independence; however, red (R), green (G) and blue (B) primaries are the most commonly used.

The expression (1.1) shows that the LMS coordinates of a color are a linear function of its spectrum, and since the RGB intensities are related to the LMS values by the linear transformation (1.3), the results of the color matching experiments should satisfy a linearity principle. Indeed, this property was discovered experimentally long before the structure of the human eye was understood, and is commonly stated in the following

> **Grassman's law:** If a test source with spectrum $I_1(\lambda)$ is matched by a triplet of primary intensities (R_1, G_1, B_1), and a second test source with spectrum $I_2(\lambda)$ is matched by the triplet (R_2, G_2, B_2), then the light obtained by mixing the two sources, which has a spectrum $I_1(\lambda) + I_2(\lambda)$, is matched by the primary triplet $(R_1 + R_2, G_1 + G_2, B_1 + B_2)$.

Or, equivalently,

$$\left.\begin{array}{l} I_1(\lambda) \stackrel{\text{match}}{=\!=\!=} (R_1, G_1, B_1) \\ I_2(\lambda) \stackrel{\text{match}}{=\!=\!=} (R_2, G_2, B_2) \end{array}\right\} \quad \Rightarrow \quad I_1(\lambda) + I_2(\lambda) \stackrel{\text{match}}{=\!=\!=} (R_1 + R_2, G_1 + G_2, B_1 + B_2) \quad (1.5)$$

As previously mentioned, a portion of the visible colors can not be matched using three predetermined primaries. However, the matching becomes possible also in this case if one or two primaries are mixed with the test source. We shall indicate with $(R_{\text{add}}, G_{\text{add}}, B_{\text{add}})$ the primary triplet added to the original source and with (R, G, B) the triplet measured in the match. Following the linearity principle expressed by Grassman's law (1.5), the primary components corresponding to the original source are computed by subtracting the contribution of the primaries mixed to the test source. The resulting triplet can therefore contain one or two negative components:

$$I(\lambda) + I_{\text{add}}(\lambda) \stackrel{\text{match}}{=\!=\!=} (R, G, B) \quad \Rightarrow \quad I(\lambda) \stackrel{\text{match}}{=\!=\!=} (R - R_{\text{add}}, G - G_{\text{add}}, B - B_{\text{add}}) \quad (1.6)$$

The color matching experiments are repeated for different monochromatic test sources spanning the entire visible spectrum. If the resulting data are plotted versus the wavelength λ of the test source, a set of curves known as *color matching functions* is obtained. By Equation (1.3), the color matching functions corresponding to an arbitrary set of three linearly independent primaries are a linear combination of the spectral sensitivities $\rho_L(\lambda)$, $\rho_M(\lambda)$ and $\rho_S(\lambda)$ of the cone cells. By their physical nature, spectral sensitivities are non-negative functions. Following these principles, in 1931 the CIE proposed to estimate the cone sensitivities by taking three different linear combinations of color matching functions which produce a non-negative value for all values of λ. The resulting combinations are indicated with the symbols $\bar{x}(\lambda)$, $\bar{y}(\lambda)$ and $\bar{z}(\lambda)$

and are tabulated; a plot is visible in Figure 1.3. As a further normalization, the transformation was chosen in such a way that the color matching function $\bar{y}(\lambda)$ corresponds to the photopic efficiency $V(\lambda)$. The color coordinates corresponding to a spectrum $I(\lambda)$ are indicated with the symbols X, Y and Z and are computed with the expression

$$X = \int \bar{x}(\lambda)\, I(\lambda)\, d\lambda \qquad Y = \int \bar{y}(\lambda)\, I(\lambda)\, d\lambda \qquad Z = \int \bar{z}(\lambda)\, I(\lambda)\, d\lambda. \tag{1.7}$$

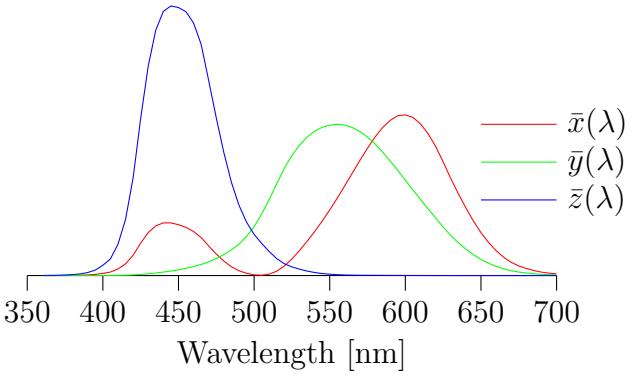

Figure 1.3: Plot of the CIE-XYZ color matching functions

The XYZ coordinates do not correspond exactly to the LMS cone response values, in particular due to the presence of the constraint $\bar{y}(\lambda) = V(\lambda)$. Different transformations have been proposed; for instance, the International Color Consortium [36] recommends the Bradford model, defined by the following matrix [42]:

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{bmatrix} 0.8951 & 0.2664 & -0.1614 \\ -0.7502 & 1.7135 & 0.0367 \\ 0.0389 & -0.0685 & 1.0296 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \tag{1.8}$$

1.6 Detection threshold and just noticeable differences

As a consequence of the retinal cell interactions described in Section 1.3, the perceived brightness of a scene is a nonlinear function of the physical luminance L. This nonlinearity is often measured psychophysically in *detection threshold* experiments. Typically, an observer looks at a uniformly illuminated screen on which some test pattern is superimposed. The pattern is visible only if its difference in luminance from the background is high enough, and the threshold value is commonly known as *just noticeable difference* (JND). If the experiment is repeated with a different background level, the measured JND changes accordingly, and by plotting the JND versus the background luminance, a *threshold versus intensity* (TVI) curve is obtained. The JND is a statistical, rather than an exact quantity: the same stimulus can be detected in some test trials and not in others, even by the same observer, and this ambiguity is most likely to happen for intensities around the threshold value we seek to measure. The JND is therefore usually defined as the stimulus intensity which is detected 50% of the times. This definition is not universally agreed upon; different researchers might use different definitions, and therefore report different results. Other properties of the human visual system can be

measured using techniques related to the concept of JNDs. The spatial sensitivity, for example, is typically measured by means of sinusoidal gratings. The minimum amplitude of the grating which is distinguishable from a uniform background depends both on the background level and on the spatial frequency; the reciprocal of the threshold, plotted versus the frequency of the grating (measured in cycles per degree of visual field) for a fixed background luminance, is commonly known as *contrast sensitivity function* [7]. Color sensitivity can also be measured, although the data is more difficult to model. For example [46], an observer can mix three primary colors in order to match a given sample. Because of the limited sensitivity of the human eye, the color found by the observer does not exactly match the reference; the JND is estimated by repeating the experiment several times and computing the error variance.

In the 19th century, the German physician Ernst Heinrich Weber discovered that the JND increases with the background luminance and that for a wide range of values the relationship is well approximated by a direct proportionality. More precisely, the relationship discovered by Weber is $\mathrm{TVI}(L) \approx k\,L$, where the constant k takes a value of approximately $0.02 \div 0.05$. This property holds for different kinds of physical stimuli and is commonly known as Weber's law. Intuitively, Weber's law can not hold for arbitrarily low luminance values, since this would imply that the human eye has an infinite sensitivity and is able to distinguish arbitrarily close luminance values as the background luminance tends to zero. Indeed, detection threshold experiments showed that Weber's law fails for luminance levels below 100 cd/m^2 – which is the range of common display devices – and the TVI curve tends to a finite asymptote as the luminance tends to zero.

Different analytical expressions for the TVI curve were proposed, depending on the test pattern used in the experiments and on the definition of threshold. In the early 1970s, Blackwell [14] suggested the following expression for the TVI curve, using a briefly flashing dot as a test target (the luminance is measured in cd/m^2):

$$\mathrm{TVI}(L) = 0.0594\,(1.219 + L^{0.4})^{2.5}. \tag{1.9}$$

With a similar experiment, Ferwerda [20] proposed the following expression:

$$\log_{10}\mathrm{TVI}(L) = \begin{cases} -0.72 & \text{if } \log_{10} L \leq -2.6 \\ \log_{10} L - 1.255 & \text{if } \log_{10} L \geq 1.9 \\ (0.249\log_{10} L + 0.65)^{2.7} - 0.72 & \text{otherwise.} \end{cases} \tag{1.10}$$

This relationship is viable for luminances approximately above 1 cd/m^2 (photopic vision); a similar expression was derived for scotopic (rod-based) vision, and some Authors [78, 1] use a combination of the two curves to describe the behavior of the HVS for all luminance ranges.

It is possible to define a mapping $L \mapsto P(L)$ such that an increase in L equal to 1 JND corresponds to a constant (typically unitary) increase in $P(L)$ [45]. This operation is commonly known as *perceptual linearization*, and the function $P(L)$ itself is called *capacity function* by some Authors [1]. From the definition above, we request that

$$P\big(L + \mathrm{TVI}(L)\big) = P(L) + 1. \tag{1.11}$$

By approximating $P\big(L + \mathrm{TVI}(L)\big)$ with its first-order Taylor expansion around L, we obtain a differential equation which can be used to compute the capacity function $P(L)$ corresponding to a given TVI function:

$$P'(L)\ \mathrm{TVI}(L) = 1 \qquad \Rightarrow \qquad P(L) = \int_{L_0}^{L} \frac{dl}{\mathrm{TVI}(l)}, \tag{1.12}$$

where L_0 is an arbitrary starting point such that $P(L_0) = 0$. For instance, if the TVI function is approximated by Weber's law, a logarithmic response is deduced:

$$P(L) = \int_{L_0}^{L} \frac{dl}{k\,l} = \frac{1}{k} \log \frac{L}{L_0}. \tag{1.13}$$

The perceptual distance between two arbitrary luminance levels L_1 and L_2 is generally assumed to be proportional to $|P(L_2) - P(L_1)|$. The function $P(L)$ represents the number of distinguishable luminance levels, or JNDs, between L_0 and L; for this reason, it is often interpreted as a subjective scale which describes the relationship between physical luminance and perceptual brightness. Many tone mapping algorithms, such as [78] and [1], are based on a curve derived from Ferwerda's model [20] of the HVS.

1.7 Display devices and gamma correction

As outlined in Section 1.2, one important consequence of the trichromatic nature of the human eye is that a wide gamut of colors can be reproduced by mixing three primary colors – typically red, green and blue (RGB) – with appropriate intensities. Following this property, most emissive display devices form an image by means of an array of pixels, each composed by three modulated light sources, or *subpixels*, of red, green and blue color. As previously noticed in Section 1.2, this representation is a consequence of one particular technology used in the reproduction of color and is not an intrinsic property of its physical and physiological nature; for this reason, the representation is not uniquely defined and different alternatives are possible which involve a different choice of the primaries. The most common standard is sRGB [72], in which the RGB intensities are computed from the XYZ values (1.7) by means of the following linear transformation:

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.2410 & -1.5374 & -0.4986 \\ -0.9692 & 1.8760 & 0.0416 \\ 0.0556 & -0.2040 & 1.0570 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \tag{1.14}$$

The format of most digital image file formats is closely linked to the display technology: each pixel in the image is represented by three numbers, typically 8 bit integers ranging from 0 to 255, which encode the intensity of the three primary subpixels. When the image is displayed, the graphics adapter converts these digital values into analog electrical waveforms which are transmitted to the display, typically through the VGA cable. For two reasons described in the following, it is advisable to use a nonlinear quantization for the RGB values encoded in an image file.

Until recent times, most display devices were based on the cathode ray tube (CRT) technology. The luminance of the phosphors in a CRT is proportional to the current of the cathode ray; this current, however, is a nonlinear function of the applied voltage and the relationship is well approximated by a power law:

$$I \propto V^{\gamma} \qquad \gamma \approx 2.2. \tag{1.15}$$

In order to display an image without distortion, the RGB intensity values should be mapped with the inverse power law. This operation is commonly known as *gamma correction* from the Greek letter γ commonly used for the exponent in the power law (1.15) [60]. In computer displays, the gamma correction is typically accomplished by mapping the digital pixel values

prior to display. These values are then converted into an analog waveform by the graphics adapter, which generates a voltage proportional to the digital values it receives. Liquid crystal displays (LCD) are based on a completely different technology and have a different response to the applied voltage; however, LCDs contain circuitry which maps the electrical signal in order to simulate the power law response of CRTs.

One second issue is related to the nonlinear behavior of the human visual system described in Section 1.6. Since the just noticeable difference increases with the luminance, the eye is able to discriminate more finely spaced levels in the dark areas than in the bright ones. Therefore, a uniform quantization would produce a poor rendition of the dark areas unless a high bit depth is used. In order to avoid this problem and to obtain a satisfactory image quality at a lower bit depth, it is convenient to use a nonuniform quantization with a spacing between the levels which follows the behavior of the threshold versus intensity function. More precisely, the linear RGB values should be mapped with an appropriate nonlinear function, which resembles the capacity function introduced in Section 1.6, and finally encoded with a uniform quantization.

By a useful coincidence, the two mappings curves – the gamma correction curve used to compensate the nonlinear voltage-to-luminance relationship in the cathode ray tubes and the capacity function used to reduce the visibility of quantization noise – have a similar shape; therefore, the gamma correction also has the effect of reducing the visibility of quantization noise, because the power law response of a cathode ray tube compresses the levels in the dark areas where the function (1.15) has a lower slope.

Since the power law response of the CRTs can be easily measured directly, whereas the capacity function must be estimated in detection threshold experiments and can vary among different observers, the former is the most used in consumer display devices. A simple power law $V \propto I^{1/\gamma}$ still does non provide good results, both because the curve has an infinite slope in the origin which can amplify the noise and give rise to numerical problems, and because the luminance of a CRT screen also contains an offset due to the reflection of ambient light. The international standard recommendation ITU-R BT.709 [37] proposes a modified curve with an offset for the compensation of the ambient light; in the lower levels where the curve would take negative values, it is replaced by a tangent linear segment:

$$D = \begin{cases} 1.099L^{0.45} - 0.099 & \text{if } L > 0.018 \\ 4.5L & \text{otherwise} \end{cases}, \tag{1.16}$$

where D is the normalized digital value and L is the normalized desired luminance. We shall see, in Chapter 5, that the display devices and the image encoding in some professional applications such as the medical one use a nonlinear encoding which is not based on the curve (1.16), but rather follows exactly an appropriate capacity function in order to minimize even further the visibility of the quantization noise and to adapt to the luminance range of each specimen. Medical display devices contain circuitry which guarantees that the output levels be spaced correctly following the desired curve within a predetermined tolerance.

1.8 Image processing pipeline in digital cameras

The CCD or CMOS sensor of a digital camera is made of an array of photodiodes. In order to discriminate color, a *color filter array* (CFA) is placed in front of the photodiodes. The most

common design, introduced in 1975 by Bryce E. Bayer of Eastman Kodak [8], uses a periodic pattern of 2×2 pixels: one sensitive to the short (S), two to the middle (M) and one to the long (L) wavelengths. This arrangement was chosen because the human eye is most sensitive to the intermediate wavelengths, which approximately correspond to green. The current generated by the photodiode during the exposure time is accumulated in a capacitor, and at the end of the exposure the voltage is amplified and read by an analog-to digital converter. Due to the presence of thermal noise in the circuits and amplifiers, the measured voltage can take negative values, especially when a high sensitivity or ISO speed is used, which involves a strong amplification of the photocurrent. For this reason, a constant offset is typically added to the digital signal in order to allow an unsigned encoding.

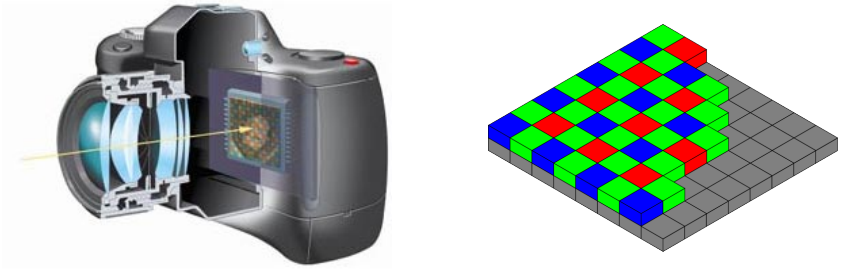

Figure 1.4: Bayer color filter array in a digital camera

In order to convert the raw data into an RGB image which can be displayed on a screen, a sequence of processing steps must be performed. The first operation, known as *demosaicing*, interpolates the color filter array in order to recover a triplet of LMS values for each pixel. Demosaicing algorithms are still an active research topic and are critical for the generation of an aesthetically pleasant image; indeed, low-quality demosaicing algorithms can significantly amplify the noise which is inevitably present in the sensor data and introduce artifacts such as aliasing in presence of high-frequency details and false colors around edges.

Once the three LMS channels have been recovered, a *white balancing* is performed in order to compensate the color of the light source. Typically, each LMS channel is multiplied by a suitable constant, with the objective that the white or gray objects in the input image appear achromatic when displayed on the screen. The white point of a display device corresponds to the brightest color the display is able to reproduce and is clearly obtained by driving the three RGB subpixels at the maximum power. Therefore, the RGB coordinates of white are $(1, 1, 1)$ in normalized units and achromatic colors have equal components. The coordinates of the white point in the LMS raw image depend on the light source, and can be estimated in different ways. In case of *preset* white balancing, the user selects the light source (daylight, tungsten, fluorescent, blackbody radiation, etc.) from a menu and the camera uses some precomputed multipliers. In case of *automatic* white balancing, the camera attempts to estimate the light source by analyzing the image data. One simple and common technique divides each LMS channel by its mean value, based on the *gray world* assumption that the spatial average of the image pixels produces an achromatic color. In case of *manual* white balancing, the user calibrates the camera by pointing it towards a white or achromatic object.

After the white balancing, the scaled pixel values are converted from the LMS color space of the sensor to the RGB color space of the display. The *color space conversion* is linear and is computed using an appropriate 3×3 matrix. The LMS color space is not uniquely defined;

an ideal camera should reproduce the spectral sensitivity of the cones in the retina, but the human photopigments are complex proteins which can not be integrated into an image sensor chip. Therefore, commercial sensors use bandpass filters which approximate the eye response, and the manufacturer provides a matrix used to convert from LMS to the standard XYZ color space. In our experiments, we used a Canon EOS 5D and a Canon EOS 350D digital reflex camera. The color transformation matrix of the EOS 5D camera is:

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{bmatrix} 0.6347 & -0.0479 & -0.0972 \\ -0.8297 & 1.5954 & 0.2480 \\ -0.1968 & 0.2131 & 0.7649 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \tag{1.17}$$

The color transformation matrix of the EOS 350D camera, instead, is:

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{bmatrix} 0.6018 & -0.0617 & -0.0965 \\ -0.8645 & 1.5881 & 0.2975 \\ -0.1530 & 0.1719 & 0.7642 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \tag{1.18}$$

The XYZ data are then converted to the RGB color space used by the display device. Finally, the RGB values are gamma corrected and converted to 8 bits.

Digital cameras perform this processing internally and produce an RGB image ready for display. The image is typically compressed in JPEG format in order to reduce its size and increase the number of shots which can be stored in one memory card. Professional cameras can also save the raw output of the sensor, thus allowing the user to perform all the processing at a later stage. In most cases, the raw image files are stored in proprietary formats which compress the sensor data using a lossless algorithm and add a certain number of metadata and tags. For instance, the white balance multipliers are simply stored in a tag and the pixel data are not modified. Although proprietary, most raw image formats are based on the TIFF specification and can be decoded using third-party software such as DCRAW.

Chapter 2

High dynamic range image composition from multiple exposures

2.1 Introduction and motivation

The first problem encountered when handling high dynamic range images is the acquisition of the image itself. In an exposed and developed film, the *optical density* – defined as the base 10 logarithm of the reciprocal of the transmittance, determined by the concentration of silver particles deposited during the development process – is a function of the *exposure* – defined as the product of the intensity of the incident light and the exposure time – following a law commonly known as the *characteristic curve* or the *Hurter-Driffield curve*. An example is shown in Figure 2.1. The curve has a sigmoidal shape and, for an intermediate range of exposure values, the optical density is a linear function of the logarithm of the exposure with a slope commonly known as *gamma*. For high exposures, the slope decreases and the curve forms a horizontal asymptote known as *shoulder*. Therefore, the amplitude of the details in the overexposed parts of a photograph gradually decreases and eventually disappears. For low exposures, the slope decreases in a similar way and the curve forms another horizontal asymptote known as *toe*. Therefore, the details in the underexposed parts of a photograph are also attenuated and eventually disappear. For this reason, conventional film is able to capture with full detail only a limited range of exposure values.

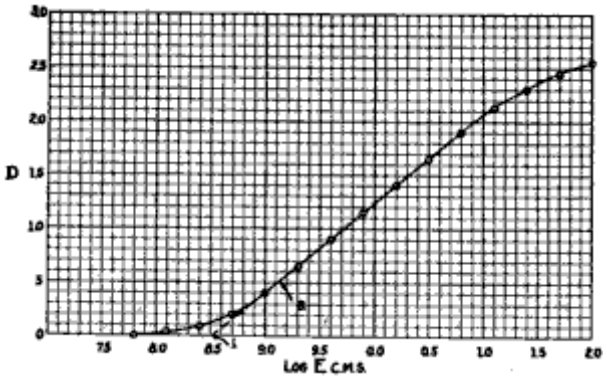

Figure 2.1: Example of density versus exposure curve of conventional photographic film

In digital photography, the range of exposure values which can be captured correctly by the sensor is also limited. The voltage generated by a CCD or CMOS is, to a good approximation, a linear function of the exposure value; however, the analog to digital conversion necessarily introduces a saturation and the pixels brighter than a certain threshold are clipped. Unlike analog film, the clipping occurs abruptly rather than gradually and the overexposed pixels are replaced by a uniform white area with a total loss of detail. Although the sensor response is theoretically linear, for low exposure values the quantum noise in the incident light, the thermal noise inside the sensor circuitry and finally the quantization noise introduced by the converter can surpass the signal intensity and compromise the visibility of the details.

Both in analog and in digital photography, it is possible to change the global brightness of the image in a wide range by adjusting the aperture, shutter speed and ISO rating, but once these parameters are set, the useful dynamic range of a single exposure is limited to about $2 \div 3$ orders of magnitude depending on the gamma of the film or on the bit depth and noise characteristics of the sensor. Therefore, if a photographer attempts to capture a high dynamic range scene, and for instance adjusts the camera settings to correctly expose the bright areas, the details in the dark areas are lost due to noise and an *underexposed* photograph is obtained such as the one in the left part of Figure 2.2. Likewise, if the photographer adjusts the camera settings to correctly expose the dark areas, the bright areas get clipped and an *overexposed* photograph is obtained such as the one in the right part of Figure 2.2.

Figure 2.2: Example of underexposed and overexposed photographs

2.2 The 1997 Debevec algorithm

One popular technique, introduced by Debevec [16], attempts to acquire a high dynamic range image by taking multiple exposures with different exposure times. This technique is still the most widely used, although it is clearly applicable only to static scenes and the original algorithm makes a simplifying assumption which is not always close to reality.

The output of a digital sensor is proportional to the product, called *exposure*, of the incident light intensity and the exposure time. This law, commonly known as *reciprocity*, holds to an excellent extent in digital sensors; reciprocity failure can occur in analog film, especially for very long exposure times, but in the following we shall concentrate on digital acquisition only. As a consequence of the reciprocity law, the luminance in an image can be estimated, up to a scaling factor, by dividing the sensor output by the exposure time. After this scaling, the

different exposures can be combined by means of an appropriate pixel-by-pixel weighted average, in which the weights privilege the correctly exposed pixels and discard the underexposed or overexposed ones. However, in many consumer digital cameras the raw sensor output is not accessible, and the files saved on the memory card, typically in JPEG format, have undergone several transformations such as white balancing, color space conversion and gamma correction. The method proposed in [16] makes the simplifying assumption that these transformations can be approximated by three independent nonlinear mappings on the three RGB channels of the image. More precisely, the digital output y of the camera for a given pixel and spectral band is assumed equal to

$$y = f(L\,T), \tag{2.1}$$

where L is the input luminance, T is the exposure time and $f(\cdot)$ is a nonlinear, monotonically increasing function to be estimated. Due to the monotonicity, the function $f(\cdot)$ is invertible; we can therefore write

$$L\,T = f^{-1}(y) \qquad \Rightarrow \qquad \log L + \log T = \log f^{-1}(y) \triangleq g(y) \tag{2.2}$$

Since the output values y are quantized, the function $g(\cdot)$ is actually a discrete look-up table. If different photographs of the same scene with different exposure times are available, it is possible to estimate the inverse camera response curve, up to a scaling factor, by means of an interpolation technique. If no measurement errors are present, the pixels in the different exposures satisfy the system of equations

$$g(y_{i,j}) = \log L_j + logT_i, \tag{2.3}$$

where the index i identifies the photograph, taken with an exposure time T_i, and the index j identifies the pixel. In order to reduce the processing time and the memory consumption, it is advisable to use only a small subset of the image pixels in the computation. The unknowns are the values of the function $g(y)$ for each digital value y and the logarithm of the luminance L of each pixel used in the interpolation.

Due to the inevitable presence of measurement errors, introduced for instance by the noise, quantization and variations between the response of different pixels, the system in Equation (2.3) should be treated in a least-squares sense. Moreover, since the darkest and brightest pixels in the image are likely saturated, a lower weight should be assigned to the corresponding equations. It is reasonable to set the weight equal to a function $w(\cdot)$ of the digital value y_{ij}. With the addition of the weighting factor, the expression to minimize becomes

$$\sum_{i=1}^{M}\sum_{j=1}^{N} w(y_{ij})\big[g(y_{ij}) - \log L_j - \log T_i\big]^2 = \text{Min} \tag{2.4}$$

The curve $g(\cdot)$ estimated by minimizing the expression (2.4) inevitably has an irregular behavior due to the presence of numerous errors; moreover, the optimization problem is likely ill-conditioned or even underdetermined, for instance when some values of y are absent. In order to overcome these issues, a regularization term is introduced in order to privilege a smooth curve. One reasonable choice for the regularization term involves the discretized integral of the square of the second derivative. This term is also weighted in order to give a lower importance to the extremal values, which are discarded anyway in the composition of the image as described in the following. If we consider 8 bit data, with the digital values y ranging from 0 to 255, the

expression to be minimized becomes

$$\sum_{i=1}^{M}\sum_{j=1}^{N} w(y_{ij})\left[g(y_{ij}) - \log x_j - \log t_i\right]^2 + \lambda \sum_{y=1}^{254} w(y)\left[g(y-1) - 2\,g(y) + g(y+1)\right]^2 = \text{Min}, \tag{2.5}$$

where the scalar λ adjusts the tradeoff between the curve smoothness and the interpolation accuracy. The expression (2.5) is still underdetermined, since it does not change if a constant value is added to g and subtracted to $\log L$. In other words, the inverse camera response curve can only be estimated up to a scaling factor. In order to obtain a unique solution, the value of $g(y)$ for one digital value y is assigned arbitrarily. One possible choice is to request that the highest output value correspond to a unitary exposure, i.e. $g(255) = 0$ since we are operating in the logarithmic domain.

After the inverse response curve $g(\cdot)$ is estimated, the different exposures are combined in order to produce a single high dynamic range image. The method proposed in [16] involves the following steps:

1. Compute the log exposure of each pixel in the source images, using the $g(\cdot)$ function estimated previously. For color images, each RGB channel is processed separately.
2. Compute the log luminance of each pixel by subtracting the logarithm of the exposure time.
3. Compute a weighted average of the images. The weight for each pixel j is a function of the digital pixel value y_{ij} and privileges the intermediate values, discarding the darkest and brightest ones which are likely saturated; it is possible to use the same function $w(\cdot)$ used in Equation (2.4).
4. Take the exponential of the assembled image

In other words, the log luminance of each pixel is estimated using the expression

$$\log \hat{L}_j = \frac{\sum_i w(y_{i,j})\left[g(y_{i,j}) - \log T_i\right]}{\sum_i w(y_{i,j})} \tag{2.6}$$

2.3 Image composition from raw sensor data

One major limitation of the method proposed by Debevec [16] and recalled in the previous Section is that the processing performed by a digital camera, as described in Section 1.8, is significantly different from a simple nonlinear mapping of the three RGB channels. This approximation can lead to the formation of artifacts in the assembled image and often results in a poor rendition of the colors. In order to overcome this limitation, it is advisable to take the exposures using a camera capable of saving raw image files. The use of raw image data has several advantages. The three LMS channels in a raw image can be considered truly independent because each one comes from a different subset of pixels in the color filter array. Therefore, the interpolation method proposed in [16] can be used to estimate the sensor response on a more solid theoretical ground. For high quality digital reflex cameras, the sensor response

is actually linear to an excellent approximation. We made some experiments using a Canon EOS 5D and a Canon EOS 350D cameras and found that for both models the sensor indeed has a linear response, therefore the estimation of the sensor response curve can be avoided altogether. The sensors in both cameras produce a 12 bit output: in the EOS 5D camera, the black level corresponds to the pixel value 128 and the white level to 3692; in the EOS 350D camera, the black level corresponds to 256 and the white level to 4095. Although not all levels are used, the resolution is higher than the one of an RGB image, which typically has 8 bits per channel produced by a nonlinear mapping which is often scarcely documented, and may be further degradated by the lossy JPEG compression which involves a decimation of the chrominance and a quantization of the DCT coefficients. Besides the higher bit depth of the data, the raw image has not yet undergone the demosaicing process which often amplifies the noise significantly. Despite these advantages, raw image files are not widely used in practice; one main inconvenience is that most camera vendors use proprietary file formats for the raw images and offer a software package for the conversion to an RGB format such as JPEG or TIFF. Recently, open source programs such as DCRAW allow to decode the files and access the raw unprocessed data.

After decoding the raw files using the DCRAW software, the images can be assembled using a technique based on a maximum-likelihood statistical estimation [63] which aims at minimizing as far as possible the noise present in the original images. The model assumes that the raw pixel value y_i produced by the camera with an exposure time T follows a Gaussian distribution with mean value $L\,T$ and variance σ^2. The variance is generally not known in advance, but its value is not critical and can be estimated heuristically. With this assumption, the probability density function (PDF) of the raw pixel value y_i has the well known expression

$$f(y) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(y - L\,T)^2}{2\sigma^2}\right). \tag{2.7}$$

By taking N photographs with different exposure times T_i, we obtain a sequence of digital values y_i. If the scene is static and the camera aperture is not changed, we can assume that L is constant in time. It is also reasonable to assume that each measurement y_i depends only on the corresponding exposure time T_i, and is not influenced by the previous operation of the camera; therefore, the measurements are statistically independent and their joint probability is equal to the product of the marginals. The joint conditional PDF of the vector $(y_1, y_2, \ldots, y_N)$ given the pixel luminance L is also known as the *likelihood function* and has the expression

$$\mathcal{L} \triangleq \prod_{i=1}^{N} \frac{1}{\sigma_i\sqrt{2\pi}} \exp\left(-\frac{(y_i - LT_i)^2}{2\sigma_i^2}\right). \tag{2.8}$$

The maximum likelihood method estimates the unknown parameters in the probability distribution of a random variable by computing the values which maximize the likelihood function corresponding to a set of observations. In the present case, we wish to estimate the pixel luminance L by maximizing the likelihood function (2.8) of the raw values captured in the multiple exposures. It is easier to maximize the logarithm of $\mathcal{L}$:

$$\log\mathcal{L} = \sum_{i=1}^{N} \left[-\log\left(\sigma_i\sqrt{2\pi}\right) - \frac{(y_i - LT_i)^2}{2\sigma_i^2}\right] = \text{Min.} \tag{2.9}$$

In order to compute the maximum, we take the derivative of (2.9) with respect to L and set it

equal to zero:

$$\frac{\partial \log \mathcal{L}}{\partial L} = \sum_{i=1}^{N} \frac{T_i \left(y_i - L\, T_i\right)}{\sigma_i^2} = 0 \qquad \Rightarrow \qquad L = \frac{\sum_i y_i T_i / \sigma_i^2}{\sum_i T_i^2 / \sigma_i^2} \tag{2.10}$$

$(y_1, y_2, \ldots, y_N)$ The pixel values y_i close to the white level are likely saturated and should not be used in the estimate. Formally, this can be accomplished by assigning a high or even infinite value to the variance σ_i^2. Ideally, it is not necessary to discard the darkest levels because their value is only affected by noise, which is handled by the maximum likelihood estimation method. A real sensor, however, may exhibit some nonlinear behavior near the extremal values or an offset due to an incorrect compensation of the dark current; if this occurs, it is advisable to discard also the pixels which take values close to the black level. In general, the noise variance σ_i^2 for each exposure can be set equal to an appropriate function of the pixel value y_i. The choice of this function is not critical in the algorithm, because the expression (2.10) is based on a weighted average of multiple observation and minor changes in the weights introduce a negligible change in the output; in particular, it can be noticed that the luminance L estimated using the expression (2.10) does not change if the variances are multiplied by the same constant.

After the assembly, the LMS images are converted to RGB using the standard procedure described in Section 1.8. The two cameras we used have a sufficiently high resolution; therefore it is optionally possible to avoid the demosaicing process altogether and to remove the color filter array by simply decimating the image. In figure 2.3 we present a false color representation of the luminance of a scene estimated with the method described in this Section. We took 17 photographs at intervals of 1 stop, with a fixed f/4 aperture and exposure times ranging from 1/8000 to 8 seconds; two of the exposures were shown previously in Figure 2.2. The values are in normalized units, with 1 corresponding to the saturation level of the sensor at an exposure time of 1 second, 100 ISO sensitivity and f/4 aperture.

Figure 2.3: False color representation of the luminance in an HDR scene estimated with multiple exposures. The values are in normalized units, with 1 corresponding to the saturation level of the sensor at an exposure time of 1 second, 100 ISO sensitivity and f/4 aperture.

Chapter 3

Dynamic range reduction of photographic images

3.1 Introduction and motivation

The intensity of natural light can span over 10 orders of magnitude, ranging from around $5 \cdot 10^{-5}$ lux in starlight to around 10^5 lux in sunlight. Even in a single scene, the *dynamic range* or *contrast* – defined as the ratio between the highest and the lowest luminance level – can easily reach values of $10^3 \div 10^4$, and sometimes exceed 10^6 in presence of critical illumination conditions involving for instance backlit objects or dim indoor scenes with strong sunlight entering from a window [20]. On the other hand, most hardcopy and display devices are able to correctly reproduce luminance values in a range of $2 \div 3$ orders of magnitude at most; therefore, the faithful reproduction of a *high dynamic range* (HDR) image is not possible in general. However, studies on the human visual system (HVS) showed that the perceived brightness of each point in a scene is not simply determined by its absolute luminance, but rather on its spatial or temporal variation [81, 39]. This remarkable property, commonly known as *local adaptation*, is actually a consequence of different types of physiological and psychophysical mechanisms: the pupil controls the global amount of light entering the eye, the bleaching and regeneration of the photopigments in the retinal receptors adjust their sensitivity on a longer time scale, and the neural interactions among the retinal cells modulate their response as a function of the local mean luminance known as *adaptation level*. As a consequence, if the luminance of an image is adjusted in a way that reduces the large-scale variation without altering the fine details, it may be possible to reduce the dynamic range without producing a significant change in the visual sensation experienced by the observer.

Following the technological progress of the image reproduction techniques, different kinds of *dynamic range reduction* methods were introduced. Even before the invention of photography, many artists had the ability of reproducing, on the limited luminance range and color gamut of a painting on canvas, the visual impression produced by the observation of a natural scene. In the example of Figure 3.1, a scene containing a very bright light source and dark unlit regions is reproduced with great naturalness. In conventional photography, it is possible to adjust the dynamic range of the printed image by carefully selecting the type of paper used for the print and by manually adjusting the exposure of different portions in the image with a technique commonly known as *dodging and burning*. Digital photography, combined with the theoretical

Figure 3.1: Example of artist's rendering of an HDR scene: "The Alchemist discovering phosphorus" by Joseph Wright of Derby. Oil on canvas, 1771.

advances in digital signal and image processing and with the continuously increasing processing power of computers, opened endless possibilities and dynamic range reduction algorithms are an active research topic. In recent times, such algorithms – once confined to professional and research applications – are being implemented in consumer imaging devices and software.

Dynamic range reduction algorithms can be divided into two main categories. *Global* operators map each pixel individually with a fixed, spatially-invariant function, or *tone mapping curve*. *Local* operators, instead, can map pixels with the same input luminance to different output levels depending on their neighborhood. Both categories have advantages and drawbacks. Global operators are usually faster and easier to tune, and the one-to-one mapping from real-world to display luminances can be a desirable property in some applications, such as the medical ones, where the luminance values of the image carry a critical information content. This same correspondence, however, inevitably reduces the luminance variations in the image and can reduce the visibility of the fine details. Local operators, on the other hand, are typically able to preserve the fine details in the image and can produce higher-quality results; however, this preservation can introduce *gradient reversal* effects because the same input value can be mapped to different output values depending on the surroundings. Indeed, simple local operators can introduce artifacts such as halos around the object boundaries, and the prevention of such artifacts can increase the computational cost significantly.

3.2 Nonlinear mapping of grayscale images

A dynamic range reduction algorithm is basically an operator which maps high dynamic range real-world input luminance values to low dynamic range output values suitable for display or print. The most simple way to achieve this effect is to map each pixel individually with a fixed, spatially-invariant function, or *tone mapping curve*. A simple linear mapping, such as gain adjustment, is not suitable for this purpose. If the luminance is attenuated in order to display the brightest pixels without saturation, the dark pixels will be covered by quantization noise and ambient light and become uniformly black. The peak brightness of a computer monitor or TV screen (approximately $100 \div 300$ cd/m^2) is also much lower than the brightness of a real-life scene, therefore, even without quantization noise and in optimal viewing conditions, the darkest pixels might fall below the sensitivity of the human eye. If instead the gain is adjusted in order to preserve the visibility of dark pixels, the bright pixels will be clipped. Indeed linear mapping, which is included in most image processing softwares with the name of *levels adjustment*, is typically used to *expand* low contrast images. It is therefore necessary to use a *nonlinear* function, which attenuates the bright pixels in order to avoid or minimize clipping, while keeping dark pixels above quantization noise and the threshold of visibility.

In a wide range of luminances, the eye sensitivity to luminance variations is approximately described by Weber's law, and following a reasoning outlined in Section 1.6, the perceptual sensation of brightness is proportional to the logarithm of the physical luminance. One simple tone mapping curve, inspired by this property, is obtained by applying a linear mapping to the *logarithm* of the luminance:

$$\log L_{\text{out}} = \gamma \log L_{\text{in}} + C \qquad \Rightarrow \qquad L_{\text{out}} \propto L_{\text{in}}^{\gamma}. \tag{3.1}$$

The expression (3.1) is commonly known as *gamma correction* because of the letter used for the exponent. Despite its simple expression, the gamma correction technique has some drawbacks which discourage its use. A power law curve with exponent $\gamma < 1$ has an infinite slope in the origin and this can cause an excessive amplification of the noise in the dark portions of the image. Moreover, the value of the exponent γ depends on the image content and a manual adjustment is necessary for each image. Finally, the theoretical motivation derived from Weber's law is weak because the law is approximately valid only for high luminances, above the white level (a few hundred cd/m^2) of most display devices.

A popular tone mapping curve, based on a modified logarithmic function, was introduced by Drago et al. [17]. One peculiarity of the proposed function is that the base of the logarithm is varied adaptively; more precisely it is set to 2 for dark pixels, 10 for bright pixels, and for intermediate values it is computed by interpolating between 2 and 10 with a "bias" power function:

$$L_{\text{out}} \propto \log_{\beta}\big(1 + L_{\text{in}}/L_0\big) \qquad \beta \triangleq 2 + 8(L_{\text{in}}/L_{\max})^{\log(b)/\log(0.5)}. \tag{3.2}$$

The Authors suggest to set the parameter L_0 equal to the logarithmic average of the input luminance and the bias parameter b equal to 0.85. Finally, the curve is normalized in such a way that the maximum value $L_{\max}$ of the input image is mapped to the display white level. With these choices, the method produces satisfactory results for most images without the need of a manual parameter adjustment.

In [27], we propose a simpler expression based on a modified logarithm with a constant base

and a single parameter adjusting the shape:

$$L_{\text{out}} \propto \log\left(1 + \frac{L_{\text{in}}}{L_0}\right). \tag{3.3}$$

Also in this case, the curve is normalized in such a way that the maximum value $L_{\max}$ of the input image is mapped to the display white level. Unlike the adaptive logarithmic mapping method by Drago et al., setting the parameter L_0 in the expression (3.3) equal to the logarithmic average of the input luminance does not provide aesthetically pleasing results. However, we have found experimentally [28] that good results are obtained by setting L_0 equal to the 25th percentile of the input luminance. In Chapter 4 we will give a more detailed description of the experimental setup and results.

One alternative approach, based on the concept of *histogram equalization*, builds a custom curve based on the statistical properties of the input image rather than using an analytical expression. The basic concept of histogram equalization is that the limited dynamic range of the display device is best exploited if the output image has an uniform histogram. This is typically accomplished by dividing the input dynamic range into uniformly spaced bins and mapping the image with a piecewise linear curve whose slope in each bin is proportional to the corresponding pixel count. This method is totally automatic, however its theoretical justification is objectionable, and indeed a simple histogram equalization has two major limitations. The first one is that, since the human visual system has a nonlinear response, the request of a uniform output histogram tends to produce unbalanced images in which the bright areas appear more heavily compressed than the dark ones[1]. The second one is that the equalization can over-enhance in some cases the portions of the image corresponding to highly populated regions of the histogram; besides producing unnatural-looking images, this reduces the available output dynamic range even further. The method introduced by Ward [78] overcomes the first limitation by equalizing the histogram in logarithmic domain, which approximately represents the perceptual brightness if Weber's law is assumed to hold; if the physical luminance values are available, the method can be further improved by using a threshold versus intensity function which takes into account the loss of sensitivity for low luminance levels. In order to overcome the second limitation, a "linear ceiling" is introduced which limits the slope of the tone mapping curve. The method is called "histogram adjustment" rather than "histogram equalization", because the output histogram is typically not uniform any more due to the constraint. The method generally produces high quality results and is still widely used despite its age.

3.3 Color processing in dynamic range reduction algorithms

Dynamic range reduction algorithms, such as the ones outlined in Section 3.2, can be applied to color images in mainly two alternative ways. One possibility is to apply the operator to each of the three color channels individually. This operation is preferably performed in the LMS color space, imitating the process that apparently occurs in the human eye and is at the base of the color constancy mechanism outlined in Section 1.3 [50]; however, digital images are most commonly encoded in the RGB color space, and this space is also used in the processing for

[1] This artifact is less visible if the histogram equalization is performed on nonlinear (gamma-corrected) data.

practical reasons. With this approach, the computational cost of the processing is increased because the dynamic range reduction algorithm must be performed three times. Moreover, this method alters in general the ratios between the color channels and can change their appearance, typically reducing the saturation. This effect can be disturbing or pleasing depending on the input image. In particular, if the input image contains a color cast, for instance due to an incorrect white balancing, a dynamic range reduction of the three RGB channels individually can reduce the cast and improve the aesthetic appearance of the image. If instead the colors in the original image are already balanced correctly, their alteration is generally unwanted.

Figure 3.2: Nonlinear mapping on luminance only, with the luminance computed as maximum (left) and weighted average (right) of the RGB channels.

One alternative approach is to extract a luminance channel Y_{in} from the input image, process it to produce an output luminance Y_{out} and then scale the three color channels proportionally [66, 75]. More precisely, the following expression is used:

$$R_{\text{out}} = R_{\text{in}} \frac{Y_{\text{out}}}{Y_{\text{in}}} \qquad G_{\text{out}} = G_{\text{in}} \frac{Y_{\text{out}}}{Y_{\text{in}}} \qquad B_{\text{out}} = B_{\text{in}} \frac{Y_{\text{out}}}{Y_{\text{in}}} \tag{3.4}$$

The luminance of a color image can be computed in different ways; the most common ones use the maximum of the three RGB channels, following the definition of the HSV color space [69], or an appropriate weighted average, such as for instance the Y component of the XYZ color space[2] which is proportional to the photopic luminance. A comparison of the two techniques, applied to the logarithmic mapping method proposed in [27], is visible in Figure 3.2. Despite the apparently weaker theoretical motivation, the first method has some practical advantages which make it preferable. The method has a lower computational cost and provides an easier control on the range of the output image, because the fractions in Equation (3.4) are smaller than 1 and each color channel of the output image is therefore upper bounded by the mapped luminance Y_{out}. We have also verified that the output images obtained with the first method are typically richer in detail and more pleasing to the observer. One possible explanation is that the three primary colors have a different perceptual brightnesses, as outlined in Section 1.4: for instance, a green light appears brighter than a blue light. If the luminance Y_{in} is computed following photometric considerations, the dynamic range reduction algorithm will tend to increase the luminance of blue areas and decrease the luminance of green areas, thus producing a more flat-looking image in which the color differences appear reduced. If instead the luminance channel Y_{in} is computed by taking the maximum of the three RGB channels,

[2]This should not be confused with the Y component of the YUV color space, which is a weighted average of *nonlinear* (gamma corrected) RGB values, and has therefore a different meaning despite the same symbol.

the different perceptual brightness of the different color hues is preserved when the luminance is mapped by the dynamic range reduction algorithm.

The image computed by (3.4) is finally gamma corrected and quantized for display. Some Authors [19] perform no gamma correction and compute the output color image with the modified expression

$$R_{\text{out}} = Y_{\text{out}} \left(\frac{R_{\text{in}}}{Y_{\text{in}}} \right)^s \qquad G_{\text{out}} = Y_{\text{out}} \left(\frac{G_{\text{in}}}{Y_{\text{in}}} \right)^s \qquad B_{\text{out}} = Y_{\text{out}} \left(\frac{B_{\text{in}}}{Y_{\text{in}}} \right)^s \tag{3.5}$$

with $s \approx 0.4 \div 0.6$. We believe that (3.4) is theoretically more valid, since the gamma correction compensates the nonlinearities of the display and should therefore be performed at the end of any image rendering algorithm.

Figure 3.3: Proposed logarithmic mapping applied to luminance only (left) and to each RGB channel individually (right).

In conclusion, processing the three color channels individually can be preferable when the input image contains an unwanted color cast; if instead the colors in the original image are correctly balanced, the dynamic range reduction algorithm should operate on the luminance only in order to preserve the ratios between the color channels. A comparison of the two methods, combined with the proposed modified logarithmic mapping [27] and applied to the "Memorial" test image, is visible in Figure 3.3. The original image has a strong red-yellow cast, which was probably introduced by an automatic white balance algorithm in order to compensate the strong blue light coming from the skylight. In this case, a mapping of the three RGB channels individually attenuates the color cast and produces a visually more pleasing image. Some methods [78] also take into account the loss of color sensitivity for low luminance values. Besides having a higher computational complexity, this approach is only possible if the physical luminance values in

the input image are known. The advantage of using this kind of methods is arguable, since they artificially degrade the image quality; their application is therefore limited to applications – such as for instance the design and simulation of lighting systems – in which the mapped image should reproduce the actual visibility of real-life scene as accurately as possible.

3.4 Retinex-based methods and the 1983 Frankle-McCann algorithm

The intensity of the light reaching a camera or an observer's eye can be described as a pointwise product of two components: the *illumination*, which depends on the light sources in the scene, and the *reflectance*, which instead depends on the surfaces of the objects:

$$L(x, y) = I(x, y) \cdot R(x, y). \tag{3.6}$$

In the 1960s and 1970s, Edwin Land and John McCann conducted a series of psychophysical experiments to study the visual appearance of objects illuminated by different light sources. The experiments, outlined in Section 1.3, suggested that the visual sensation mostly depends on the reflectance and that the human visual system is able to discard the effect of illumination sources with both a nonuniform spectrum and spatial distribution. This processing apparently takes place partially in the retina and partially in the visual cortex of the brain; for this reason, the model is generally known under the name of *Retinex* as a contraction of the words *retina* and *cortex*. The Authors attempted to design a computational algorithm which simulates this behavior of the human visual system in order to estimate the reflectance from a captured digital image [44, 43]. The method initially proposed estimated the reflectance of each pixel in the image by constructing a certain number of random paths across the image leading to the pixel under consideration. The processing makes use of an appropriate sequence of operations commonly known as *ratio*, *product*, *reset* and *average*. As a first starting guess, the reflectance of the farthest pixel in each path is initialized to 1; the reflectance of the subsequent pixels in the path is estimated as follows. First, the ratio is computed between the luminance values of the current and the previous pixel in the path. The result of the ratio operation only depends on the reflectance and is insensitive to the illumination as long as the illumination is sufficiently smooth. The reflectance of the current pixel is then estimated by multiplying the estimated reflectance of the previous pixel and the ratio. If the result is greater than 1, the reflectance is reset to 1 following the assumption that the objects reflect only a fraction of the incident light. This process is repeated for all the pixels of the path and for all the paths leading to the pixel under consideration. In this way, a number of estimates for the pixel reflectance equal to the number of random paths is obtained. Finally, these values are averaged together to produce the output. Although this method attempts to imitate the behavior of the eye, which continuously makes small saccadic movements when fixating an objects, the results of the algorithm have a poor quality and suffer from a high amount of random noise. Moreover, the method contains a very high number of parameters, such as the number, length and shape of the paths, which influence the output and are difficult to tune.

In 1983, Frankle and McCann [21] patented a variation of the algorithm which uses a single and deterministic path with a spiral shape rather than multiple random paths. Besides simplifying the computation, the use of a fixed path avoids the formation of random noise. The algorithm was also designed for implementation on a specialized hardware (the International Imaging

Systems I²S image processor) and exploits some specific capabilities of that hardware, in particular the shifting of image planes.

The algorithm operates on the logarithm of the input image, in order to replace multiplications and divisions with simpler additions and subtractions. We indicate with $\mathcal{L}(x, y)$ the logarithmic luminance of the input image, and allocate a temporary image $OP(x, y)$ known as "old product", which is initialized to 0. The algorithm then proceeds as follows. The ratio and product steps produce a temporary image $IP(x, y)$ known as "intermediate product", which is computed as

$$IP(x, y) \triangleq OP(x + \Delta x, y + \Delta y) + \mathcal{L}(x, y) - \mathcal{L}(x + \Delta x, y + \Delta y),$$

where Δx and Δy are the displacements defined by the spiral pattern which will be described in the following. The reset step takes the intermediate products which are greater than zero and replaces them with 0. The result, indicated with $IP^*(x, y)$, is computed as

$$IP^*(x, y) \triangleq \min\big\{IP(x, y), 0\big\}$$

Finally, the average step computes a "new product" $NP(x, y)$ by averaging the reset intermediate product $IP^*(x, y)$ with the old product:

$$NP(x, y) \triangleq \frac{IP^*(x, y) + OP(x, y)}{2}$$

After these four steps, the old product is replaced by the new product and the process is iterated with the next displacement values Δx and Δy.

The spiral path is made of a sequence of *revolutions*. In each revolution, every pixel is compared with the one d steps to the right ($\Delta x = d$, $\Delta y = 0$), then above ($\Delta x = 0$, $\Delta y = d$), then left ($\Delta x = -d$, $\Delta y = 0$) and finally below ($\Delta x = 0$, $\Delta y = -d$). Each revolution is repeated a certain number of times, then the displacement d is divided by 2 and the processing is repeated until d becomes equal to 1. The initial value of d is typically chosen equal to the largest power of 2 which is smaller than the image width and height.

When the algorithm terminates, the new product contains the logarithm of the estimated reflectance in the scene. Due to the reset step, the reflectance is guaranteed to be less than 0 in logarithmic units, or 1 in linear units, in agreement with the observation that the objects reflect in general only a fraction of the incident light. The output is displayed by taking the exponential and performing a gamma correction.

3.5 Mathematical analysis and improvements of the Frankle-McCann algorithm

It is interesting to make a mathematical analysis of the Frankle-McCann algorithm. The motivation is twofold. First, with the original formulation it may be difficult to understand and visualize the effect of the different steps. By rewriting the method in a different but equivalent way, also used in [70], the operations become more clear. In particular, we will show that the algorithm is able to compute a nontrivial operation at a remarkably low computational cost. Second, it is easier to make some modifications on this alternative formulation, in order to further reduce the computational cost, reduce the formation of artifacts and adapt the operator to different applications which will be described in the following chapters.

As a first step, we merge the ratio, product, reset and average operations into a single expression, eliminating the IP and IP^* variables. Since the old product and the new product converge to the logarithm of the estimated reflectance, we rename them to $\mathcal{R}$ and $\hat{\mathcal{R}}$ respectively. For simplicity of notation, we also define $\mathbf{p} \triangleq (x, y)$ and $\mathbf{q} \triangleq (x + \Delta x, y + \Delta y)$. At the first step, the reflectance is initialized to 0. The result of one iteration can be written as

$$\hat{\mathcal{R}}(\mathbf{p}) = \frac{\mathcal{R}(\mathbf{p}) + \min\big\{\mathcal{R}(\mathbf{q}) + \mathcal{L}(\mathbf{p}) - \mathcal{L}(\mathbf{q}), 0\big\}}{2} \tag{3.7}$$

The expression (3.7) can be modified in order to estimate the log illumination $\mathcal{I}$, rather than the log reflectance $\mathcal{R}$. By definition, in logarithmic units we have $\mathcal{L} = \mathcal{I} + \mathcal{R}$, or equivalently $\mathcal{R} = \mathcal{L} - \mathcal{I}$. At the first step, $\mathcal{I}$ is initialized to $\mathcal{L}$. By substituting this identity into (3.7), we obtain

$$\mathcal{L}(\mathbf{p}) - \hat{\mathcal{I}}(\mathbf{p}) = \frac{\mathcal{L}(\mathbf{p}) - \mathcal{I}(\mathbf{p}) + \min\big\{\mathcal{L}(\mathbf{q}) - \mathcal{I}(\mathbf{q}) + \mathcal{L}(\mathbf{p}) - \mathcal{L}(\mathbf{q}), 0\big\}}{2} \tag{3.8}$$

or, rearranging and simplifying:

$$\hat{\mathcal{I}}(\mathbf{p}) = \frac{\mathcal{L}(\mathbf{p}) + \mathcal{I}(\mathbf{p}) - \min\big\{\mathcal{L}(\mathbf{p}) - \mathcal{I}(\mathbf{q}), 0\big\}}{2} \tag{3.9}$$

By using the identity $\min(a, b) = -\max(-a, -b)$, we can rewrite Equation (3.9) as

$$\hat{\mathcal{I}}(\mathbf{p}) = \frac{\mathcal{L}(\mathbf{p}) + \mathcal{I}(\mathbf{p}) + \max\big\{\mathcal{I}(\mathbf{q}) - \mathcal{L}(\mathbf{p}), 0\big\}}{2} \tag{3.10}$$

Finally, by moving all the terms inside the max operator and simplifying, we obtain

$$\hat{\mathcal{I}}(\mathbf{p}) = \max\left\{\frac{\mathcal{I}(\mathbf{p}) + \mathcal{I}(\mathbf{q})}{2}, \frac{\mathcal{L}(\mathbf{p}) + \mathcal{I}(\mathbf{p})}{2}\right\}. \tag{3.11}$$

In order to understand the algorithm operation, we temporarily drop the max operation and only consider the first term in the curly braces. This simplified algorithm performs a sequence of convolutions with large but very sparse kernels, each consisting of two nonzero terms. We further simplify the problem by considering a one-dimensional signal and one revolution per level. The first revolution, with an initial displacement of d samples, averages each sample with the one d steps to the right, and then with the one d steps to the left. This pair of averages is repeated $\log_2 d + 1$ times, halving the displacement each time until it becomes equal to 1. The impulse response of the resulting filter, plotted in Figure 3.4, is a triangular function with a width of $4d - 1$ samples. The overall computational cost is $2(\log_2 d + 1)$ additions and an equal number of multiplications by $1/2$, which can be computed efficiently by using shift operations[3].

A full revolution in two dimensions performs one separable convolution with two identical triangular kernels, one in the horizontal and one in the vertical direction. If n revolutions are performed, the kernel is convolved with itself n times and its shape becomes that of a uniform B-spline basis function of degree $2n - 1$. A plot is shown in Figure 3.5. It can be noticed that the shape of the kernel quickly approaches a Gaussian as n increases.

[3]If a floating-point representation is used, multiplications by a power of 2 can be computed by adding an integer constant to the exponent. The Intel processor architecture has a FSCALE machine instruction which performs this task; however, this instruction has a significantly longer latency than a general floating-point multiplication and is therefore useless in the present application.

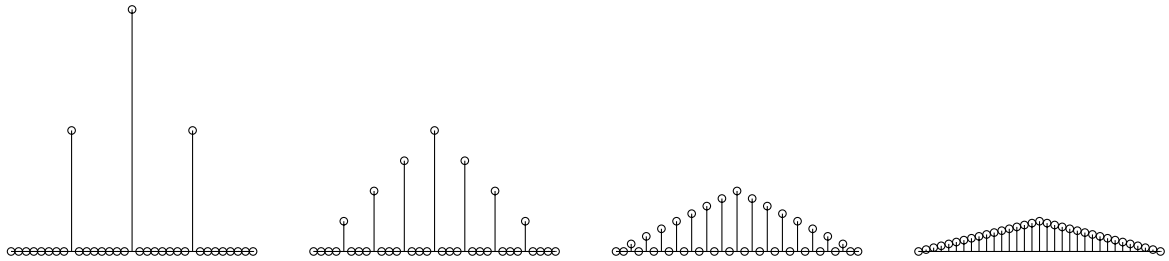

Figure 3.4: Equivalent convolution kernel of the Frankle-McCann filter with an initial displacement d of 8 samples, shown after 1, 2, 3 and 4 levels.

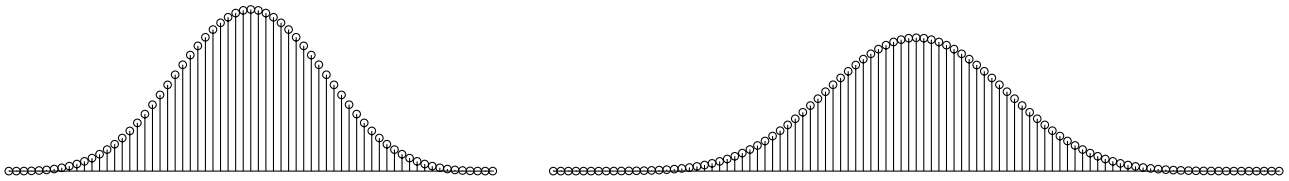

Figure 3.5: Equivalent convolution kernel of the Frankle-McCann filter with an initial displacement d of 8 samples, with 2 and 3 revolutions per level.

In conclusion, the Frankle-McCann algorithm without the reset operation performs a lowpass filter with a large, approximately Gaussian kernel at a remarkably low computational cost and using very simple operations. The memory consumption is also very low, because the processing can be performed in place without the need of allocating temporary arrays. The limitation is that, since the initial displacement d must be a power of 2 and the number of revolutions n is integer, the filter bandwidth can only take a limited set of possible values.

The reset operation has the effect of transforming the linear lowpass filter into a *constrained* lowpass filter, in which the output $\mathcal{I}$ always takes values greater than the input $\mathcal{L}$. One drawback of the original formulation is that the filter has an asymmetrical behavior. The solution we propose is to perform one single reset after each complete revolution, rather than after each iteration. We recall that one revolution is made of 4 iterations, in which each pixel in the estimated illumination is averaged with the one d steps to the right, top, left and bottom. The reset operation we propose is simply

$$\mathcal{L}^*(\mathbf{p}) = \max\big\{\mathcal{L}(\mathbf{p}), \mathcal{I}(\mathbf{p})\big\}$$

Besides producing a symmetric output, this modification provides some computational savings, because fewer reset operations are performed; it is also possible to replace the average operations by simple additions and compute one single division by 16 after one complete revolution. A comparison of the results, obtained with the original and with the modified algorithm on a simple synthetic image, is shown in Figure 3.6.

The image boundary can be handled in different ways. For instance, the implementation described in [22] pads the images with zeros; however, this method is not advisable in our opinion, because it can produce halos near the image boundaries. The method we propose is to simply skip, in each iteration, the portion of the image where the displaced pixel falls outside the canvas.

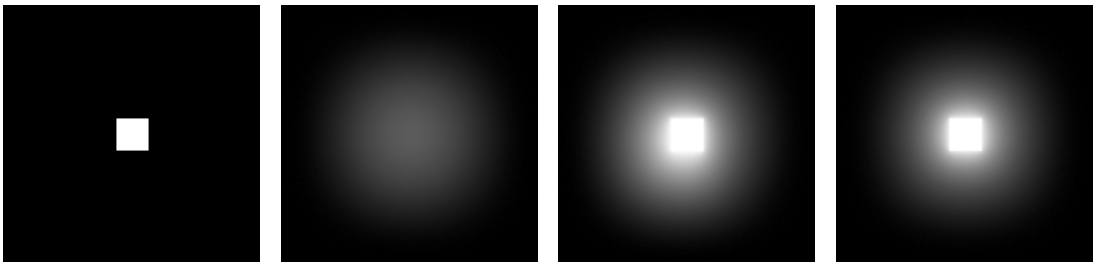

Figure 3.6: Influence of the reset operation. Left to right: input image, filter output with no reset (approximately equivalent to a Gaussian convolution), filter output with one reset per iteration (original formulation, asymmetric output) and with one reset per revolution (proposed formulation, symmetric output). The 256×256 image was processed with 2 revolutions per level and an initial displacement of 32 pixels.

It is interesting to visualize the algorithm operation by displaying the successive approximations to the reflectance after each revolution. An example, on the Memorial test image, is shown in Figure 3.7. The operator is applied to three RGB channels, and this effectively removes the reddish cast which was present in the original image. As previously noticed, a more perceptually motivated method should process the image in the LMS rather than the RGB color space; however, most test images are only available in RGB and the conversion to LMS is difficult to perform without access to the original raw exposures. The estimated reflectance has a lower dynamic range than the input luminance and is always less than 1 thanks to the reset step, which guarantees that no pixels are clipped. However, the processed image still contains some visible artifacts, in particular halos around the high-contrast edges and a bluish cast on the ceiling in the left part of the image.

Figure 3.7: Intermediate results of the modified Frankle-McCann Retinex algorithm. The 768×512 image was processed with 4 revolutions per level and an initial displacement of 256 pixels. The 9 images were visualized after each group of 4 revolutions, i.e. with a displacement of 256, 128, 64, 32, 16, 8, 4, 2 and 1 pixels.

3.6 The 1999 McCann algorithm

In 1999, McCann [50] proposed an alternative formulation of the algorithm based on a multiscale pyramidal decomposition. An approximate solution is first computed by downsampling the input image by a factor 2, processing it and upsampling the result back to the original resolution. This image is then refined by performing a certain number of Frankle-McCann revolutions with a displacement of 1 pixel. The downsampled image is processed recursively in the same way, until an image of a few pixels is obtained; at the lowest level, the log reflectance is initialized to 0.

The 1999 McCann implementation has a lower computational cost than the 1983 Frankle-McCann implementation, because part of the processing is performed at a lower resolution, although we have found that a higher number of revolutions is required in order to obtain similar results. For instance, the results of Figure 3.7, obtained with 4 revolutions per level, are similar to those obtained with the 1999 McCann algorithm with 10 revolutions.

The modified reset we have proposed above is applicable also to the multiscale implementation. In particular, the effect of one revolution is to smooth the image with the following 3×3 separable lowpass filter:

$$H = \frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

The reset operation corrects the pixels which fall below the input, thus transforming the lowpass filter into a constrained lowpass filter.

3.7 A more general formulation

The algorithms described in the previous sections attempt to separate the illumination and reflectance components in an image; after the splitting, the illumination is discarded completely and only the reflectance is displayed. Based on the assumption that the illumination is spatially smooth and that the reflectance is lower than unity, the illumination is estimated by means of a constrained lowpass filter, for which two efficient implementations were designed. This model suffers from two major drawbacks. The first one is that, in reality, the illumination is not smooth and can contain sharp edges, typically around the boundaries of objects which partially cover a light source or a brighter background. If this property is not taken into account, the estimated reflectance presents some halos around the aforementioned boundaries, as can be seen in Figure 3.7. The second one is that the human eye is not totally insensitive to the illumination, therefore discarding this component completely can produce unnatural looking images in some cases.

In order to address these issues, a more general architecture has been proposed [40]. The illumination component is estimated by means of an appropriate filter, and the reflectance is computed by division. Rather than completely discarding the illumination, this component is mapped with a suitable nonlinear function in order to reduce its dynamic range. The reflectance, instead, is generally left unaltered; some implementations [49] optionally perform some processing also on the reflectance in order to emphasize the details and to reduce the noise. Finally, the two components are multiplied together to produce the output image. A

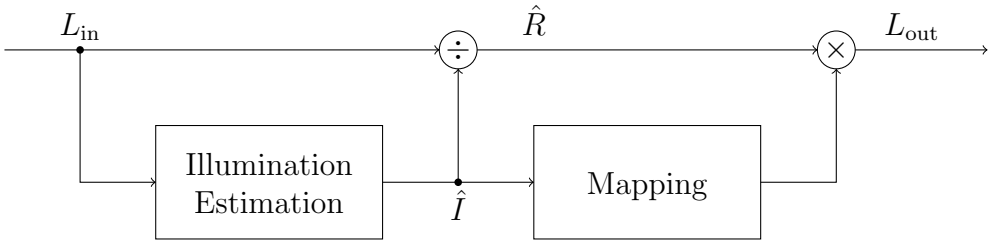

Figure 3.8: Block diagram of a general Retinex-based algorithm

block diagram of the system is shown in Figure 3.8.

The critical block in the algorithm is clearly the estimation of the illumination component. A simple lowpass filter, used for instance in early methods such as the homomorphic filter [53, 71], can introduce halos around sharp edges; a common solution to this problem is to use an *edge-preserving* lowpass filter which smooths out the low-contrast details while preserving the high-contrast edges. Several filters of this kind have been proposed in the literature, most notably in the field of noise reduction. One example is the bilateral filter, introduced by Tomasi and Manduchi [73]. Durand et al. proposed to use this filter for the estimation of the illumination component in an image (although the Authors use a different terminology) and developed two approximation techniques which reduce the computational cost [18, 54]. We also proposed a technique, based on similar principles, which allows a further acceleration [29]. Methods based on the concept of anisotropic diffusion [59] compute an edge-preserving lowpass filter by simulating the diffusion of a fluid across a porous substrate with a space-varying diffusion coefficient; the Rational Filter [61] implements this concept by using an iterative operator based on a rational function. The Recursive Rational Filter [49] is able to achieve an edge-preserving lowpass effect at a remarkably low computational cost. One example of an image processed using the bilateral filter is shown in Figure 3.9; it can be noticed that no halos are produced around the sharp edges.

Figure 3.9: Example of an image processed using the bilateral filter.

A different class of methods transforms the image into another space – such as gradient domain [34, 19] or a visual response space [48] – in which the relatively complex tone mapping task is

reduced to a more simple operation such as scaling. These methods are typically free from halo artifacts, but have a high computational cost, because the inverse transform is accomplished by means of a least-squares optimization procedure. Some methods belonging to this class also suffer from an excessive enhancement of the details which can lead to unnatural-looking images. In Figure 3.10 we show the results of the methods by Fattal [19] and Mantiuk [48], generated with the "pfstools" software package.

Figure 3.10: Example of an image processed with the methods by Fattal (left) and Mantiuk (right).

3.8 Proposed method for a Retinex-like spatially variant mapping

The bilateral filter is able to prevent the formation of halos due to its edge-preserving capability. However, the reflectance computed using the scheme of Figure 3.8 can take values above 1 in general because no constraint is present. Although the perceptual motivation of this constraint is disputed, enforcing it has some practical benefits. On grayscale images, when the mapped illumination is multiplied by the reflectance to produce the final output, the pixel values are only decreased; this automatically guarantees that no pixel will be clipped provided that the range of the nonlinear function used to map the illumination is contained in the luminance range of the output device. This property also holds for color images if the luminance Y_{in} is taken equal to the maximum of the three RGB channels, as used in the definition of the HSV color space [69]. We also noticed in our experiments that, in most cases, the reflectance computed without constraints has an unnatural appearance and can not be displayed alone; a portion of the illumination must be restored in order to obtain a visually pleasing image. In order to avoid the formation of halos and to prevent clipping, a *constrained edge-preserving* lowpass filter should be used [41, 68]. We propose an operator which achieves this task at a low computational cost and with a very limited number of user-adjustable parameters.

One consequence of the multiplicative model of Equation (3.6) is that the edges in the bright areas of the image involve a greater luminance variation compared to the edges in the dark areas, because a variation in the reflectance is multiplied by a larger value of illumination. In general, it is reasonable to assume that the edge magnitude is proportional to the local mean luminance, or in other words, that the edges are characterized by a large luminance *ratio* rather than a large luminance *difference*. The design of an edge-preserving lowpass filter is simplified

if the filter is applied to the logarithm of the luminance. In this way, the product of Equation (3.6) becomes a sum and the edge magnitude becomes uniform in the entire luminance range. Following the notation introduced in Section 3.4, we shall indicate the logarithmic variables with a calligraphic font.

One suitable design technique which allows to define the objectives and the constraints with a great flexibility is to formulate the filter as a constrained optimization problem. If the input log luminance $\mathcal{L}(x,y)$ and the estimated log illumination $\mathcal{I}(x,y)$ are treated as functions of continuous variables, one possible objective function is

$$J \triangleq \iint \left\{ w(x,y) \left\| \nabla \mathcal{I}(x,y) \right\|^2 + \underbrace{\left[\mathcal{L}(x,y) - \mathcal{I}(x,y) \right]^2}_{\mathcal{R}(x,y)} \right\} dx\, dy = \text{Min} \tag{3.12}$$

subject to the constraint

$$\mathcal{I}(x,y) \geq \mathcal{L}(x,y) \qquad \forall\ (x,y) \tag{3.13}$$

The first term in the integral privileges a smooth illumination, and the second privileges a reflectance with a low dynamic range. The space-varying coefficient $w(x,y)$ is introduced to achieve an edge-preserving effect. Ideally, this coefficient should take a large value in the smooth portions of the image, in order to privilege the lowpass effect, and a smaller value near the high-contrast edges, in order to preserve them as desired. One possible solution is to set $w(x,y)$ equal to a decreasing function of the norm of the luminance gradient, i.e. $w(x,y) = G(\|\nabla \mathcal{L}(x,y)\|)$. We tested different options, and found that good results can be obtained by choosing a coefficient $w(x,y)$ inversely proportional to the gradient norm:

$$w(x,y) \triangleq \frac{\alpha}{\|\nabla \mathcal{L}(x,y)\|}, \tag{3.14}$$

where the constant α adjusts the tradeoff between the two objectives. With this choice, α is the only parameter in the algorithm.

In order to solve the optimization problem, the objective function must be discretized. The second term is naturally replaced by a summation over the image pixels:

$$\iint \left[\mathcal{L}(x,y) - \mathcal{I}(x,y) \right]^2 dx\, dy \quad \mapsto \quad \sum_{i=1}^{M} \sum_{j=1}^{N} (\mathcal{L}_{i,j} - \mathcal{I}_{i,j})^2, \tag{3.15}$$

where M and N are the number of rows and columns in the image. The gradient term, instead, requires greater attention. If the weight factor $w(x,y)$ were not present, the most intuitive technique would perhaps be to replace the partial derivatives in the gradient with horizontal and vertical finite differences and the integral with a summation. This approach gives rise to the following expression:

$$\begin{aligned} \iint \left\| \nabla \mathcal{I}(x,y) \right\| dx\, dy = \iint \left\{ \left(\frac{\partial \mathcal{I}(x,y)}{\partial x} \right)^2 + \left(\frac{\partial \mathcal{I}(x,y)}{\partial y} \right)^2 \right\} dx\, dy \mapsto \\ \mapsto \sum_{i=1}^{M} \sum_{j=1}^{N-1} (\mathcal{I}_{i,j+1} - \mathcal{I}_{i,j}) + \sum_{i=1}^{M-1} \sum_{j=1}^{N} (\mathcal{I}_{i+1,j} - \mathcal{I}_{i,j}). \end{aligned} \tag{3.16}$$

However, if the weight factor $w(x,y)$ is introduced, the discretization is not as straightforward. If for instance each term in the summation (3.16) is multiplied by $w_{i,j}$, the resulting expression

assumes an asymmetric shape which is not desirable in general. One approach, proposed in [13], obtains a symmetric operator by summing for each pixel the squared differences in the right, left, up and down direction. We use an alternative discretization scheme which has some advantages in terms of computational cost. We consider a cell of 2×2 pixels and approximate the continuous function $\mathcal{I}(x, y)$ with a bilinear polynomial which interpolates the 4 pixels at the corners of the cell. We then compute the square norm of the gradient of this function and sample it at the *center* of the cell.

$$\begin{array}{ccc} \mathcal{I}_{i,j+1} & \cdots & \mathcal{I}_{i+1,j+1} \\ \vdots & \odot & \vdots \\ \mathcal{I}_{i,j} & \cdots & \mathcal{I}_{i+1,j} \end{array}$$

After carrying out the described computation, we obtain the following simple expression, which resembles the Roberts gradient operator [62] up to a scaling factor:

$$\|\nabla \mathcal{I}\|^2 \quad \mapsto \quad \frac{1}{2}\big[(\mathcal{I}_{i,j} - \mathcal{I}_{i+1,j+1})^2 + (\mathcal{I}_{i,j+1} - \mathcal{I}_{i+1,j})^2\big] \tag{3.17}$$

The gradient of the input luminance $\mathcal{L}(x, y)$, used in the definition of the w coefficient (3.14), is computed in the same way; therefore, also the function $w(x, y)$ is sampled at the center of the 2×2 cell. We shall indicate this value with the notation $w_{i+\frac{1}{2},j+\frac{1}{2}}$. Finally, the first term in the integral (3.12) is replaced by a summation of all the 2×2 pixel cells in the image:

$$\begin{gathered} \iint w(x, y) \left\|\nabla \mathcal{I}(x, y)\right\|^2 dx\, dy \quad \mapsto \\ \mapsto \frac{1}{2} \sum_{i=1}^{M-1} \sum_{j=1}^{N-1} w_{i+\frac{1}{2},j+\frac{1}{2}} \left[(\mathcal{I}_{i,j} - \mathcal{I}_{i+1,j+1})^2 + (\mathcal{I}_{i,j+1} - \mathcal{I}_{i+1,j})^2\right] \end{gathered} \tag{3.18}$$

One advantage of this scheme is that the expression (3.17) is consistently defined also at the image boundaries; therefore, no special treatment of the boundary conditions is required. The fact that the discretized forms of the first and second term in the objective function (3.12) are sampled in different locations is not an issue, because the samples are summed together and are not treated individually.

After the objective function is discretized, an iterative method is used to compute the minimizing point. One simple technique is the *point-by-point minimization*, in which the function is minimized with respect to one variable at a time. If we take the partial derivative with respect to one generic sample $\mathcal{I}_{i,j}$ and set it equal to zero, we obtain the equation

$$\begin{gathered} w_{i-\frac{1}{2},j-\frac{1}{2}} (\mathcal{I}_{i,j} - \mathcal{I}_{i-1,j-1}) + w_{i-\frac{1}{2},j+\frac{1}{2}} (\mathcal{I}_{i,j} - \mathcal{I}_{i-1,j+1}) + \\ + w_{i+\frac{1}{2},j-\frac{1}{2}} (\mathcal{I}_{i,j} - \mathcal{I}_{i+1,j-1}) + w_{i+\frac{1}{2},j+\frac{1}{2}} (\mathcal{I}_{i,j} - \mathcal{I}_{i+1,j+1}) + (\mathcal{I}_{i,j} - \mathcal{L}_{i,j}) = 0 \end{gathered} \tag{3.19}$$

We then solve the equation for $\mathcal{I}_{i,j}$ and update the illumination pixel with this new value. If the computed value is lower than $\mathcal{L}_{i,j}$, we replace it with $\mathcal{L}_{i,j}$ in order to satisfy the constraint. The constraint, therefore, increases only minimally the overall computational complexity. One useful property, which is a consequence of the proposed discretization scheme, is that the expression (3.19) does not contain other pixels in the same row or column of $\mathcal{I}_{i,j}$; therefore, it is theoretically possible to update all the pixels in a row or column simultaneously and the algorithm is intrinsically suitable for a parallel implementation. Modern processors are able

to compute 4 floating-point operations simultaneously using the SSE instruction set, provided that the software is appropriately coded using these instructions; the pipeline also introduces some level of parallelism if the instructions do not contain dependencies.

The point-by-point minimization scheme still requires a large number of iterations, because each step only processes pixels in a small neighborhood; in other words, the low frequencies of the image are reconstructed very slowly (although it should be noticed that the concept of frequency response can not be applied to nonlinear filters such as the present one). The convergence speed can be dramatically improved by using a multiscale method similar to the one used in the 1999 McCann algorithm described above. An approximate solution is first generated by downsampling the image by a factor 2, processing it and upsampling the result back to the original resolution. If this image is used as a starting guess, the number of iterations required for convergence is significantly smaller, because the low frequencies in the original image correspond to higher frequencies in the downsampled images. The low-resolution image is processed recursively in the same way until its size reduces to a few pixels; at the lowest level the illumination is simply initialized with the input luminance. The use of the point-by-point minimization scheme guarantees that the value of the objective function decreases at each iteration, but we have not yet developed a formal proof that the proposed method finally converges to the global minimum. Moreover, although the objective function (3.12) is a positive definitive quadratic form, the constraint (3.13) introduces local minima. However, the proposed method was able to produce a satisfactory solution for all the test images we used in our experiments.

An example of processing, obtained with the proposed method on the "window" test image used in Chapter 2, is shown in Figure 3.11. We compare the results obtained with two different values of the parameter α in order to show its effect on the processed result. As previously noticed, α is the only parameter in the algorithm and adjusts the "bandwidth" of the lowpass filter, with higher values of α producing a stronger lowpass effect in the illumination and consequently a reflectance image richer in detail. Following the scheme of Figure 3.8, we mapped the illumination with the modified logarithmic function (3.3) proposed in [27]. The dynamic range reduction algorithm is computed on the luminance only, computed as the maximum of the three RGB channels; this choice, combined with the constraint (3.13), guarantees that no pixel is clipped.

Figure 3.11: Results obtained with the proposed Retinex-like method, using $\alpha = 100$ (top) and $\alpha = 20$ (bottom) in the variational filter for the extraction of the illumination component. The estimated illumination and reflectance are shown together with the processed output. The illumination is mapped with the proposed modified logarithmic curve.

Chapter 4

Quality assessment of dynamic range reduction algorithms

4.1 Introduction and motivation

Dynamic range reduction algorithms, such as the ones described in Chapter 3, are used to process a digital image in order to reduce its dynamic range without compromising the details and the information content. By exploiting some properties of the human visual system, these methods attempt to reproduce, on a display device with a limited luminance range, a visual sensation as close as possible to the one experienced by the direct observation of the real scene.

Recently, research is being made in order to evaluate the quality of the results, compare different methods and tune their parameters. However, the definition of an objective quality metric is likely an ill-posed problem, because by its own nature a tone mapping algorithm alters significantly the luminance values in an image. For this reason, most image fidelity metrics – ranging from simple ones, such as the mean squared error, which only consider the pointwise luminance differences, to complex ones [15, 47] which include models of the human visual system – do not provide meaningful results; some metrics were proposed which only compare the local variations and are not influenced by differences in the local mean [35, 2], but their use is still at an early experimental stage. For this reason, it is often preferable to perform a subjective assessment by means of appropriate psychophysical experiments in which a group of observers are asked to rate some specific properties of the processed image or to adjust some parameters in the tone mapping algorithm.

The research we present in this chapter is targeted specifically to the quality assessment of the global tone mapping methods outlined in Section 3.2. The methods belonging to this class have a low computational cost and are easy to tune but can cause a loss of detail in the image. In Section 4.2 we introduce a *detail attenuation parameter* α_{det} which gives a quantitative measure of this artifact. The tone mapping curve proposed in [27], based on a simple logarithmic function, is shown to have a good performance in the sense that the detail attenuation is generally confined to a small portion of the image; the method is also easy to use because the expression of the tone mapping curve contains one single user-adjustable parameter. We conducted a psychophysical experiment in which several users adjusted this parameter manually in the same viewing conditions for a set of test images. The experiment is described in detail

in Section 4.3. We checked if different users agree on the same or reasonably close values of the parameter for the same test image. We also investigated how the user-adjusted parameter is related to the statistical properties of the image, and found that a percentile provides the best fit. The experimental results, described in Section 4.4, suggest a simple rule which allows to compute the parameter automatically with satisfactory results for most images.

4.2 Detail attenuation metric for global tone mapping methods

The most simple category of dynamic range reduction algorithms maps each pixel individually in the same way with an appropriate nonlinear function commonly known as *tone mapping curve*:

$$L_{\text{out}}(x,y) = F\big(L_{\text{in}}(x,y)\big). \tag{4.1}$$

The methods of this kind generally have a low computational cost and are intrinsically free from a class of artifacts, such as gradient reversals or noise amplification, which may affect operators that perform a space-dependent processing. A global mapping is also a typical building block of more advanced methods. However, global methods inevitably introduce a loss of visibility in the image details because a wide range of input luminances is mapped to a smaller, and generally finite, number of output levels. In [27] we introduce an operator which gives a quantitative measure of this artifact.

As outlined in Section 1.6, the sensitivity of the human visual system to luminance variations can be quantified by using the concept of *just noticeable difference* (JND). The JND increases with the background luminance following a law commonly known as *threshold versus intensity* (TVI) function [20], and for high luminance levels the relationship is well approximated by a direct proportionality known as Weber's law. Despite being an approximation, Weber's law is very useful in practice. One particular advantage is that the expressions derived using Weber's law as a model are often scale-invariant and do not depend on the units of measure of the physical quantity under consideration. This property is particularly desirable in image processing applications for two reasons: the pixel values of an image are generally expressed in normalized units and the corresponding physical luminance is not known; moreover, the luminance of a displayed image depends on the characteristics of the display device and is generally not known in advance. The calibration of a camera or a display device requires appropriate instruments and skills which are not accessible to the average user.

Let L_{in} and $L_{\text{in}} + \delta L_{\text{in}}$ indicate two close luminance levels in the input image. The two levels are mapped by the nonlinear function $F(\cdot)$ to output levels $F(L_{\text{in}}) \triangleq L_{\text{out}}$ and $F(L_{\text{in}} + \delta L_{\text{in}}) \triangleq L_{\text{out}} + \delta L_{\text{out}}$. Since δL_{in} is small by hypothesis, δL_{out} can be approximated by a first order Taylor expansion of F around L_{in}:

$$\delta L_{\text{out}} \triangleq F(L_{\text{in}} + \delta L_{\text{in}}) - F(L_{\text{in}}) \approx F'(L_{\text{in}})\, \delta L_{\text{in}}. \tag{4.2}$$

As outlined in Section 1.6, the perceptual distance between two close luminance levels can be estimated by taking the ratio between the luminance difference and the just noticeable difference. If we use Weber's law as a model for the TVI function, we obtain the expressions

$$\delta P_{\text{in}} \triangleq \frac{\delta L_{\text{in}}}{\text{TVI}(L_{\text{in}})} \approx \frac{\delta L_{\text{in}}}{k\, L_{\text{in}}} \qquad \delta P_{\text{out}} \triangleq \frac{\delta L_{\text{out}}}{\text{TVI}(L_{\text{out}})} \approx \frac{\delta L_{\text{out}}}{k\, L_{\text{out}}}. \tag{4.3}$$

We define a "detail attenuation parameter" α_{det} as

$$\alpha_{\text{det}} \triangleq \frac{\delta P_{\text{in}}}{\delta P_{\text{out}}} = \frac{\delta L_{\text{in}} / \text{TVI}(Lin)}{\delta L_{\text{out}} / \text{TVI}(L_{\text{out}})} \approx \frac{\delta L_{\text{in}} / L_{\text{in}}}{\delta L_{\text{out}} / L_{\text{out}}}. \tag{4.4}$$

By replacing L_{out} and δL_{out} in (4.4) with their expressions (4.1) and (4.2), we obtain

$$\alpha_{\text{det}} \triangleq \frac{\delta L_{\text{in}} / L_{\text{in}}}{\delta L_{\text{out}} / L_{\text{out}}} = \frac{\delta L_{\text{in}} / L_{\text{in}}}{F'(L_{\text{in}})\, \delta L_{\text{in}} / F(L_{\text{in}})} = \frac{F(L_{\text{in}})}{L_{\text{in}}\, F'(L_{\text{in}})}. \tag{4.5}$$

As previously noticed, the use of Weber's law allows to obtain an expression which is invariant to scaling factors in the input luminance L_{in} or the tone mapping curve $F(\cdot)$.

The derived expression measures the factor by which the image details are attenuated by the tone mapping curve, and is a function of L_{in}, since the attenuation typically depends on the original luminance. If for instance a gamma correction $L_{\text{out}} \propto L_{\text{in}}^{\gamma}$ is used, we obtain $\alpha_{\text{det}} = 1/\gamma$, thus indicating that the attenuation is uniform throughout the dynamic range. The modified logarithmic mapping proposed in [27] uses the expression

$$L_{\text{out}} \propto \log\left(1 + \frac{L_{\text{in}}}{L_0}\right). \tag{4.6}$$

For this operator, we obtain

$$\alpha_{\text{det}} = \left(1 + \frac{L_0}{L}\right) \log\left(1 + \frac{L}{L_0}\right). \tag{4.7}$$

Figure 4.1 shows a false color representation of the detail attenuation α_{det} introduced by the modified logarithmic mapping with a manual adjustment of the parameter L_0. The figure shows that the mapping preserves the image details in the dark areas and attenuates them in the bright areas. Most high dynamic range images have a strongly asymmetric histogram, with most pixels in the dark areas and with small bright areas typically corresponding to the light sources; it can also be noticed that the objects in an image generally have a greater importance than the light sources. On images with these characteristics, the detail attenuation introduced by the logarithmic mapping is confined to a small portion of the image and is scarcely visible.

4.3 Psychophysical evaluation and tuning of the logarithmic mapping operator

The tone mapping curve proposed in [27] is able to produce satisfactory results at a low computational cost, but its expression (4.6) contains a parameter L_0 which requires some sort of adjustment. It is natural to ask whether an "optimal" value of the parameter exists for a given image and a given display device, or in other words, if different users agree on the same or reasonably close values when tuning the parameter manually. If such a value exists, it may be possible to investigate how it is related to the statistical properties of the image. We conducted a psychophysical experiment [28] for this purpose.

In order to perform the experiment, we designed a graphical user interface (GUI) which displays a set of HDR images mapped using the proposed tone mapping curve (4.6) and allows the user

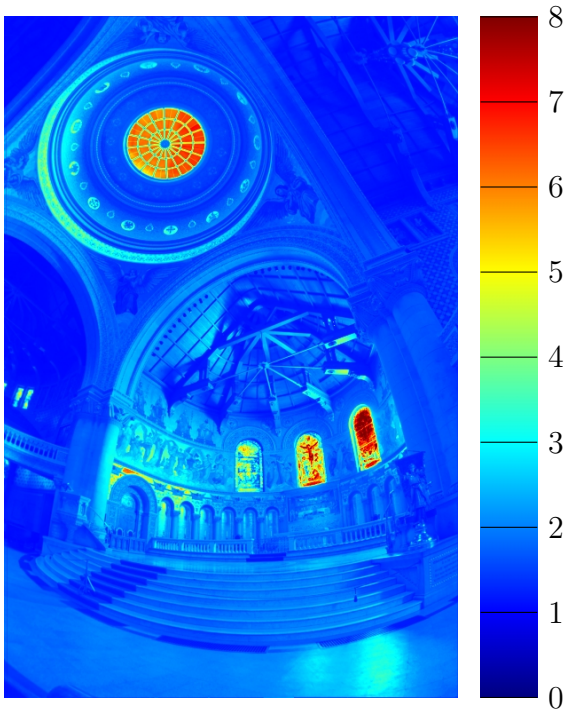

Figure 4.1: False color representation of α_{det} for the proposed modified logarithmic mapping.

to adjust the parameter L_0 interactively. When the user is satisfied with the result, the value of L_0 is saved into a log file and the next image is displayed. We used a Hewlett Packard P1100 CRT display with a resolution of 1280×960 pixels, adjusted with minimum brightness and maximum contrast[1], and repeated the experiment both with dim (a few tens lux) and with standard (a few hundred lux) ambient lighting. The experiment was repeated with different users, using a data set of 33 HDR images. No preprocessing was applied to the images, apart from downsampling them when necessary to make them fit into the display. The running time for each test trial was approximately 15 minutes.

For each image, the luminance is computed by taking the maximum of the three RGB channels. The luminance is mapped using the modified logarithmic function (4.6), and the proportionality coefficient is adjusted so that the maximum luminance L_{max} in the input image is mapped to the maximum display luminance, which we arbitrarily set to 1 in normalized units. The complete expression of the tone mapping curve is therefore

$$L_{\text{out}} = \frac{\log(1 + L_{\text{in}}/L_0)}{\log(1 + L_{\text{max}}/L_0)}. \tag{4.8}$$

The user can adjust the parameter L_0 interactively using the keyboard. Since the value of L_0 can span several orders of magnitude, we designed a simple method which allows the users to explore a large range of values in a small number of trials. At the first iteration, the GUI computes the value of L_0 which produces a tone mapping curve (4.8) with a unit slope in

[1]The brightness and contrast controls on a display are often confusing. The brightness control adds an offset to the output luminance, and should be set to zero in order to have the lowest possible black level. The contrast control varies the gain.

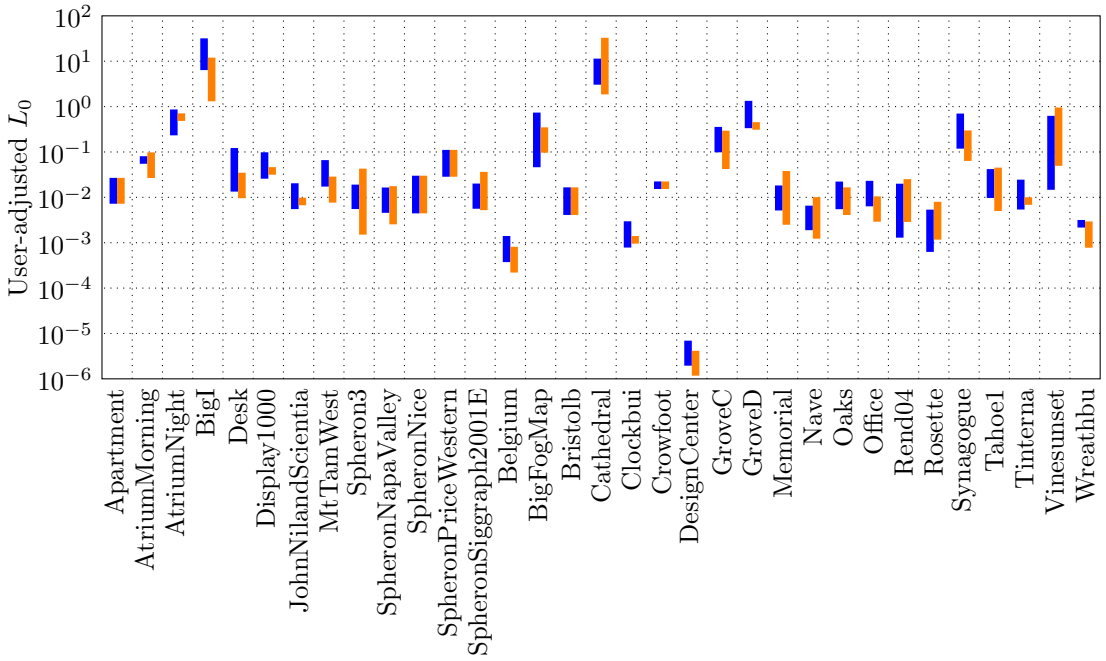

Figure 4.2: Mean and standard deviation of the parameter L_0 chosen by different users, with dim (blue) and standard (orange) ambient lighting

the origin. By pressing the + and - keys, the slope in the origin is multiplied or divided by 2 respectively, and the corresponding value of L_0 is computed by solving a simple nonlinear equation, since the normalization factor in (4.8) also depends on L_0. In other words, at each iteration the brightness of the dark areas is increased or decreased by one stop. We believe that this method allows the users to better understand and visualize the practical effect of the parameter they are adjusting, and indeed the participants were able to quickly find the preferred image even without knowing the algorithm details. We chose a step size of one stop because we found that it provides a good tradeoff between accuracy and speed.

The output color image is computed by scaling the three RGB channels proportionally using (3.4). Finally, the image is gamma corrected using the curve defined in the international standard recommendation ITU-R BT.709 [37] and quantized to 8 bits with dithering.

Thanks to the simplicity of the operator, at each iteration the image is processed and displayed in a fraction of a second. Some of the standard HDR test images included in our data set have strongly distorted colors; since the purpose of this experiment is to evaluate luminance mapping, and not color rendition, our GUI allows the users to disable color and display a grayscale-only image if they believe that color artifacts can be distracting.

4.4 Experimental results

Figure 4.2 shows an error bar plot of the parameter L_0 chosen by the different test subjects for each test image (indicated in the abscissa), with dim and standard ambient lighting. In the hypothesis that the data have a log-normal distribution, we computed the mean and the

standard deviation in a logarithmic domain. As previously noticed, the user-adjusted value of L_0 heavily depends on the characteristics of the input image; moreover, many of the HDR test images we used are encoded in arbitrary units, which are linearly proportional to the physical luminance but with an unknown scale factor. Indeed our experiments showed that the mean values of L_0 for *different* images span several orders of magnitude, therefore the error bars in Figure 4.2 naturally have different vertical positions. The values of L_0 chosen by different users for the *same* image, instead, were generally in good agreement. This can be easily measured by computing the sample standard deviation, which is represented by the length of the bars in the plot. For most test images, we obtained a standard deviation of approximately 0.5 log units. Our experiments also showed that, for most images, the effect of ambient lighting has little influence on the experimental results.

On a few test images, the agreement between different users was poor. A possible explanation is that, as noticed in Section 4.2, the modified logarithmic mapping (4.6) has a best performance on images which mostly contain dark areas. A few test images, for example outdoor scenes containing large portions of sky, do not match this hypothesis. In these cases the detail attenuation introduced by the logarithmic mapping becomes more visible, and the users are forced to set a tradeoff which is likely to depend on personal preferences. An example is the "Vine sunset" image, on which indeed we found the worst agreement.

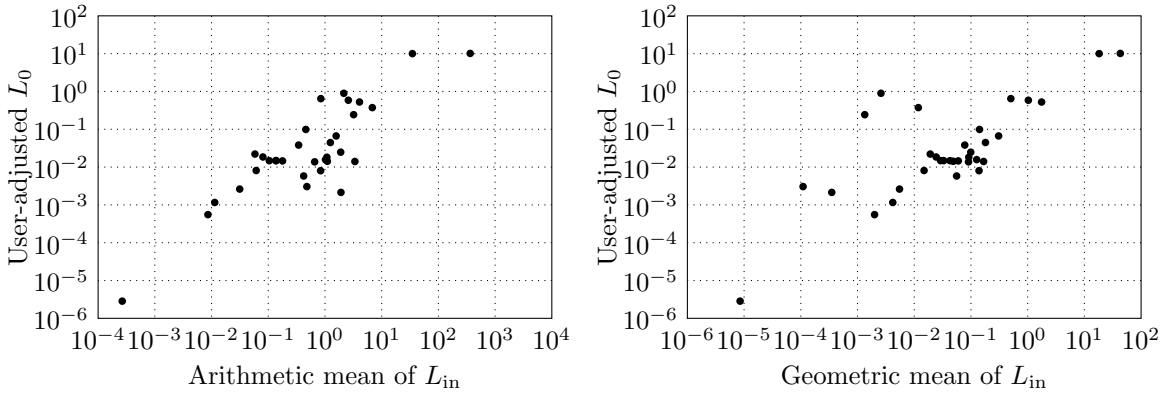

Figure 4.3: Scatter plot of the user-adjusted parameter L_0 versus the arithmetic mean (left) and geometric mean (right) of the input luminance.

In any case, since a good agreement was generally found between different users, we were able to proceed with our analysis of the experimental data by investigating if there is some simple relationship between the manually-adjusted value of L_0 and some statistical property of the image. If a good fit exists, it may be possible to automatically process an image with results comparable to those obtained by a manual tuning. It is reasonable to assume that the value of L_0 is proportional to, or at least increasing with, some sort of mean luminance value. The options we tested are the arithmetic mean, the geometric mean (or logarithmic average), and percentiles. In Figure 4.3, the values of L_0 chosen by one of the test subjects are plotted versus the arithmetic mean and the logarithmic average of the input luminance. Some trend is visible, but in both cases the fit is quite poor and there are several outliers. The sample correlation coefficients, computed in a log-log scale for the data shown in the plot, are 0.85 and 0.95 respectively. The data of the other test runs exhibit a similar behavior. The arithmetic or geometric mean are therefore unsuitable for the automatic tuning of the parameter L_0. Better results are obtained using percentiles. In Figure 4.4, the values of L_0 are plotted versus the 20th

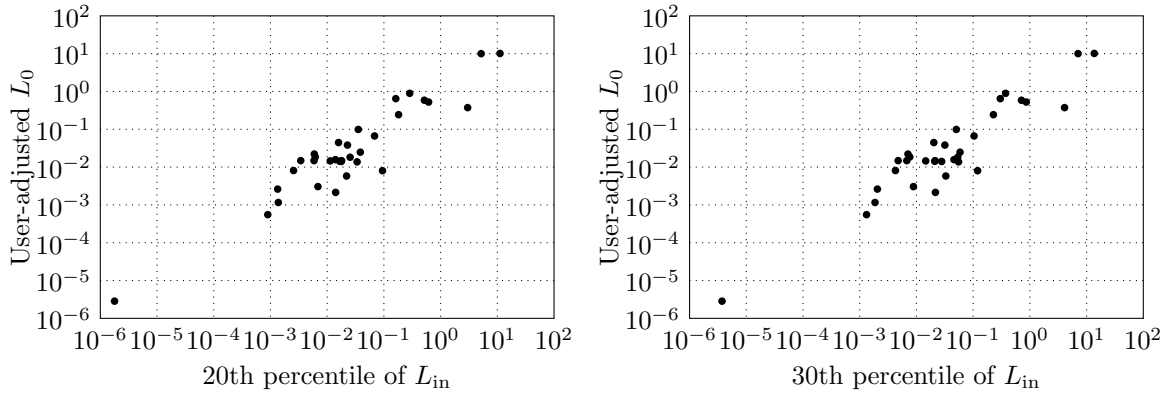

Figure 4.4: Scatter plot of the user-adjusted parameter L_0 versus the 20th percentile (left) and 30th percentile (right) of the input luminance.

and the 30th percentile of the input luminance. In this case a stronger correlation is visible, and the error which is still present in the fit has a magnitude comparable to the dispersion in the experimental data. The sample correlation coefficient is 0.93 in both cases; although this value is slightly lower than the one obtained with the logarithmic average, the plot does not exhibit the outliers which affected the former case. We have verified that the best fit is obtained by setting L_0 equal to the 25th percentile of the input luminance. A thumbnail of all the test images used in our experiments, tone mapped using the proposed method, is shown in Figure 4.5.

4.5 Conclusions and further work

In this chapter we presented two techniques for the quality assessment of dynamic range reduction methods based on global tone mapping curves. The first technique provides a quantitative measure of the detail attenuation introduced by the mapping. The modified logarithmic mapping proposed in [27] was shown to produce good results, in the sense that the detail attenuation is generally confined to a small portion of the image. The second technique uses a psychophysical experiment in order to evaluate the agreement between different users in the adjustment of the parameter L_0 in the expression of the tone mapping curve. The experiments showed that the user-adjusted parameter is, to a good approximation, equal to the 25th percentile of the input luminance. This allows to compute the filter automatically, with results comparable to those obtained with a manual tuning on most images.

The experiments described in this chapter are still in progress. We expect to perform a larger number of test trials, possibly using a calibrated display in a controlled environment. In particular, the dependence of L_0 on the characteristics of the display device and viewing environment should be further investigated. We believe that the preliminary results presented in this thesis are encouraging.

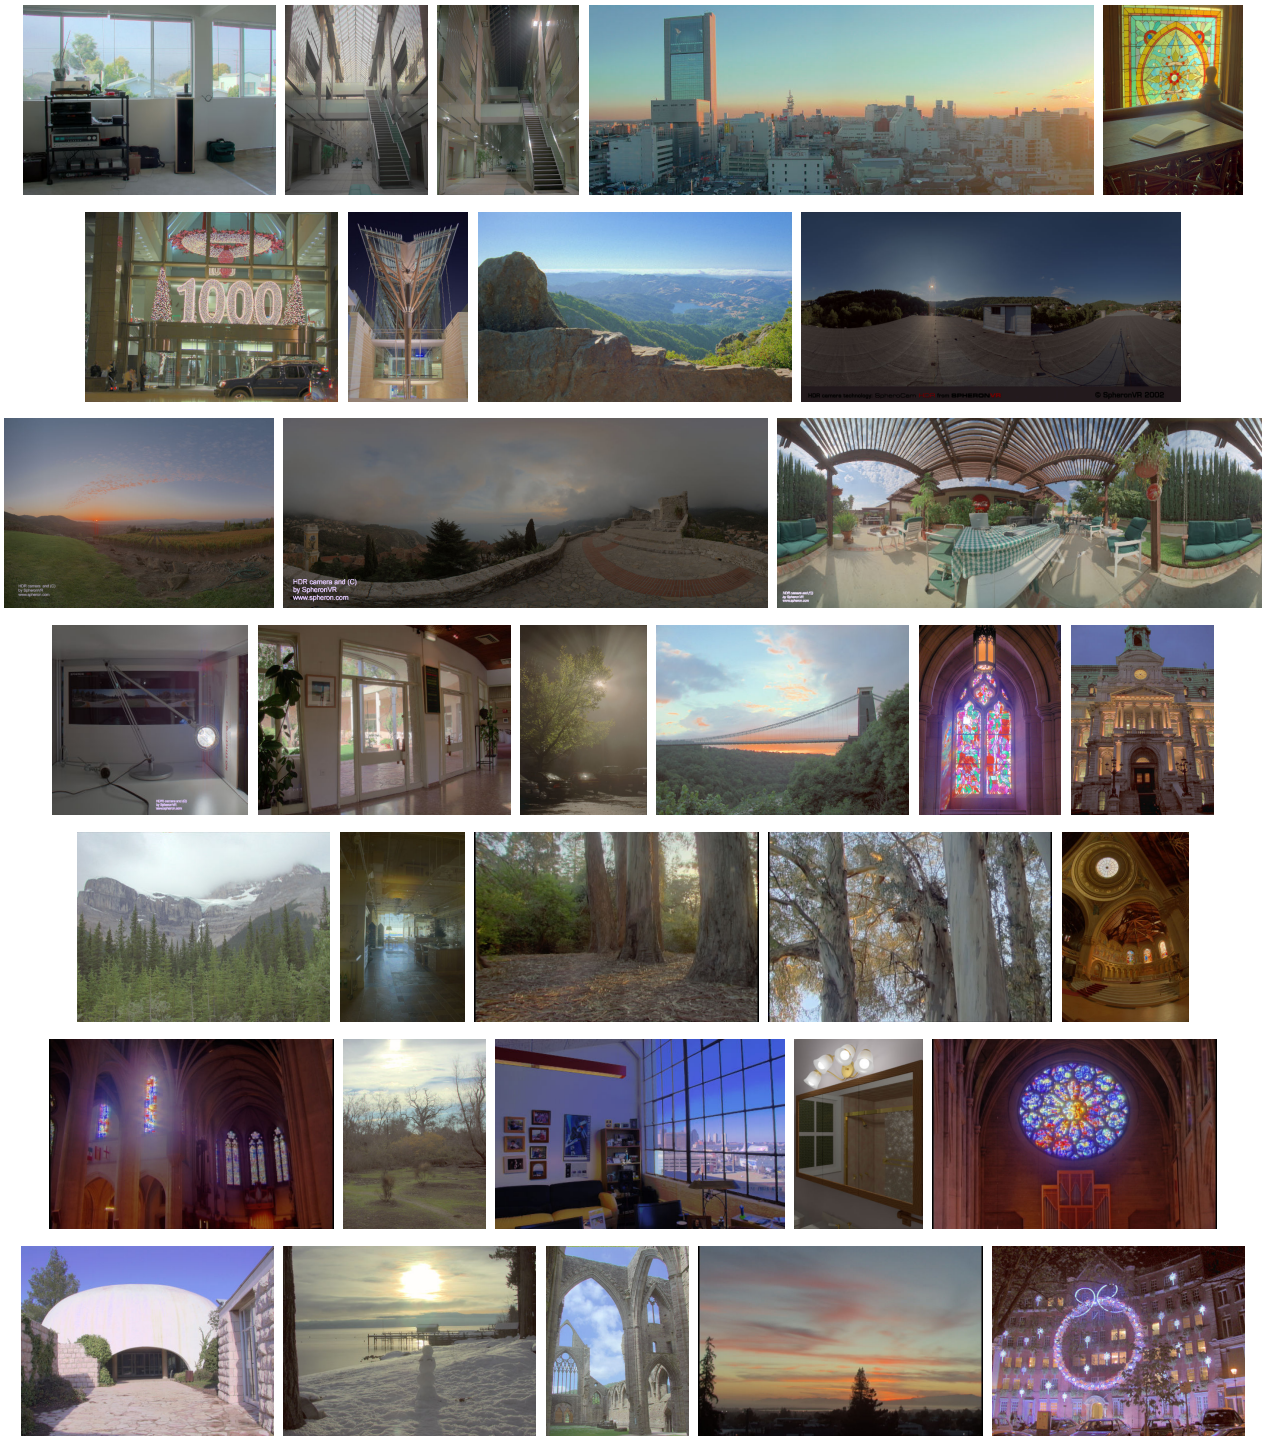

Figure 4.5: Test images used in our experiments. Top to bottom and left to right: "Apartment", "AtriumMorning", "AtriumNight", "BigI", "Desk", "Display1000", "JohnNilandScientia", "MtTamWest", "Spheron3", "SpheronNapaValley", "SpheronNice", "SpheronPriceWestern", "SpheronSiggraph2001E", "Belgium", "BigFogMap", "Bristolb", "Cathedral", "Clockbui", "Crowfoot", "DesignCenter", "GroveC", "GroveD", "Memorial", "Nave", "Oaks", "Office", "Rend04", "Rosette", "Synagogue", "Tahoe1", "Tinterna", "Vinesunset", "Wreathbu".

Chapter 5

Nonlinear mapping of radiographic images

5.1 Introduction and motivation

As outlined previously in Section 1.6, the human visual system (HVS) has a nonlinear response to luminance. This nonlinearity can be quantified using the concept of *just noticeable difference* (JND), which represents the minimum amplitude of a specified test pattern an average observer can discern from a uniform background. The JND is typically measured psychophysically in *detection threshold* experiments, in which an observer repeatedly looks at a uniformly illuminated screen on which a test pattern of varying intensity is superimposed and reports whether he is able to discriminate the pattern or not. Detection threshold experiments have been conducted repeatedly by different researchers, using different types of test patterns and experimental procedures [20]. Although the results are often significantly different, all experiments show that the JND increases with the background luminance. For high luminance levels, above approximately 100 cd/m^2, the relationship is well approximated by a direct proportionality commonly known as Weber's law. For low luminance levels, the JND increases and tends to a constant value as the background luminance tends to zero. For intermediate values, the relationship is sometimes approximated by a square root commonly known as the DeVries-Rose law. In general, the relationship between the background luminance and the JND is described by a curve commonly known as the *threshold versus intensity* (TVI) function.

This behavior should be taken into account in the encoding and visualization of digital images. If a uniform quantization of the gray levels is used, the steps are more visible in the dark areas, where the JND takes a smaller value and the human eye has a greater ability in discriminating fine luminance differences. In order to avoid this issue and to minimize the visibility of the quantization noise introduced by a limited bit depth, the digital levels should be unevenly spaced. The *gamma correction* technique, used in the encoding of conventional digital images [72, 37], approximately achieves this goal by mapping the luminance values in an image using a power law. This mapping was originally introduced to compensate the nonlinear voltage-to-luminance response of cathode ray tubes, but also has the effect of reducing the spacing between the quantized luminance levels in the dark areas, thus improving the rendition of digital images quantized at a bit depth of 8 bits per pixel. Gamma correction curves, however, are based on technological rather than psychophysical concerns and only provide a rough approximation of

the nonlinear behavior of the HVS. In professional applications, such as the medical ones, it may be convenient to choose a mapping which follows more closely the average behavior of the human visual system and to adjust the response of the display device accordingly with some circuitry which guarantees that each digital driving level (DDL) produces the desired luminance within a prescribed tolerance.

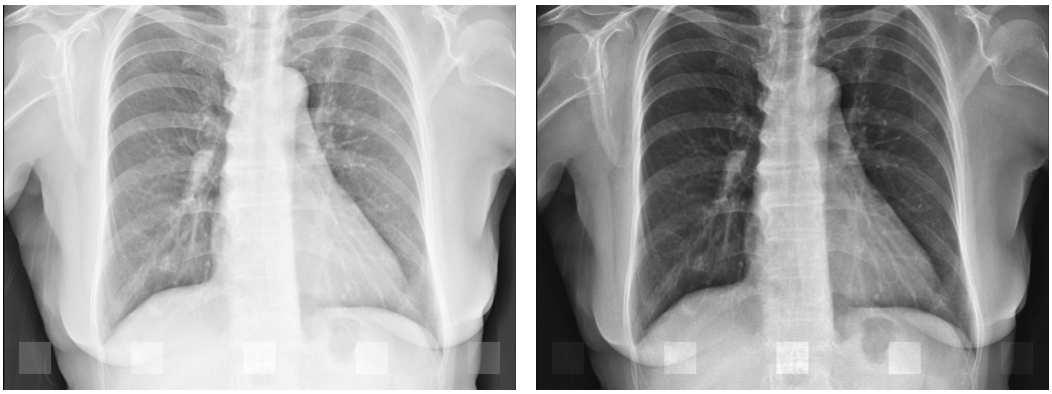

Figure 5.1: Sample radiographic image, with artificial details of the same density, mapped with a linear (left) and a DICOM (right) display function. The DICOM GSDF is calibrated for a luminance range between 1 and 200 cd/m^2.

In the specific case of the display of radiographic images, a nonlinear mapping of the luminance also serves a different purpose. An x-ray detector estimates the density of the patient's tissues by measuring the attenuation of an x-ray beam. A magnetic resonance machine measures some properties of the tissues – usually their water concentration – by means of magnetic and radiofrequency fields. In any case, the image displayed to the radiologist represents in the form of luminance some data which originally had a different physical dimension; therefore, the visualization intrinsically requires some sort of mapping. In conventional film-based radiography, the mapping is determined by the characteristic curve of the film; digital radiography, instead, uses a software look-up table which can be adjusted with a greater flexibility. It is desirable to use a mapping which equalizes the visibility of details at all luminance levels, possibly adapting to the brightness and dynamic range of the individual display device. Ideally, a well designed mapping should guarantee that objects with the same density have the same visibility in the image regardless their background. A linear mapping does not achieve this goal because the human eye has a greater ability to discriminate fine luminance differences in the dark areas than in the bright ones; therefore, if an image is rendered with a luminance proportional to the density, the objects on a dark background will appear with a greater intensity. This effect is visible in the left part of Figure 5.1: on the bottom of the image, 5 squares of equal density are superimposed, but the squares appear more visible in the parts where the background is dark. The *Digital Imaging and Communications in Medicine* (DICOM) standard recommendation, developed by the National Electrical Manufacturers Association (NEMA), introduced a mapping curve known as the *DICOM Grayscale Standard Display Function* (GSDF) which aims at equalizing the visibility of the image details in the entire luminance range [52]. The same image, mapped with the DICOM GSDF, is visible in the right part of Figure 5.1. If the image is displayed on a calibrated monitor, the squares should theoretically have the same visibility.

Despite its widespread use, the DICOM GSDF has some limitations. The analytical expression

of the curve, based on a polynomial, is applicable only for luminance values between 0.05 and 4000 cd/m^2 and diverges outside this interval. Moreover, the model used to compute the TVI function is inaccurate for very low luminance levels, and this inaccuracy can reduce the visibility of the details in these areas. High dynamic range (HDR) displays, which are beginning to appear on the market [67, 74], can display luminance levels outside the range in which the DICOM GSDF is defined; in particular, dual-layer LCD displays [76, 23] are able to reproduce very low luminance levels in which the DICOM GSDF produces a poor rendition or is undefined at all. The visibility of the details in very dark areas is further reduced by the effect of ambient light which can have significantly greater luminance than the black level of an HDR display. In this chapter, we recall the underlying theory used in the definition of the DICOM GSDF and propose a modification which attempts to extend the validity of the curve for low luminance levels and to compensate the effect of ambient light. We also propose a simple computational method for the generation of the modified curve. Unfortunately, the results can only be appreciated by viewing the mapped images on a calibrated high dynamic range display and can not be reproduced on a conventional display or on print; therefore, we only present the results by plotting the mapping curves generated using the proposed method.

5.2 The DICOM Grayscale Standard Display Function

A consequence of the nonuniform distribution of the just noticeable differences, introduced in Section 1.6 and recalled above, is that the perceptual distance ΔP between two luminance levels L_1 and L_2 is not simply proportional to their difference, but also depends on the luminance itself. If the two luminance levels L_1 and L_2 are sufficiently close, one possible method of quantifying their perceptual distance is to take the ratio between the luminance difference and the JND. Since the JND varies between L_1 and L_2, the mean value should be used; a possible expression is for instance:

$$\Delta P(L_1, L_2) \triangleq \frac{2\,(L_2 - L_1)}{\mathrm{TVI}(L_1) + \mathrm{TVI}(L_2)}. \tag{5.1}$$

Given a sequence of monotonically increasing luminance levels $L_1, L_2, \ldots, L_n$, it is possible to estimate the "perceptual brightness" P_j of each level L_j by accumulating the perceptual distances, computed using the expression (5.1), between each pair of adjacent levels in the scale. If we arbitrarily set $P_1 = 0$, the described operation becomes:

$$P_1 = 0 \qquad P_{j+1} = P_j + \Delta P(L_j, L_{j+1}) \tag{5.2}$$

The perceptual brightness P_j represents the number of JNDs, or distinguishable luminance levels, between L_1 and L_j. A luminance scale is *perceptually uniform* if the perceptual distance between each pair of adjacent levels L_j and L_{j+1} is constant, or in other words, if the luminance difference is a constant multiple of the JND. Following the definition (5.1), the desired property is

$$L_{j+1} - L_j = \frac{k}{2}\left[\mathrm{TVI}(L_j) + \mathrm{TVI}(L_{j+1})\right] \qquad \forall j. \tag{5.3}$$

A perceptually uniform luminance scale satisfying the property (5.3) can serve both of the purposes outlined above. If the quantized luminance levels of a display device are adjusted following the property (5.3), the quantization noise is distributed uniformly across the entire

luminance range and its visibility is minimized. If the density values of a radiographic image are mapped to luminance following a perceptually uniform scale, such that the index j in the scale is a linear function of the tissue density in the radiographic data, details with the same density should theoretically produce the same variation in perceptual brightness on any background level.

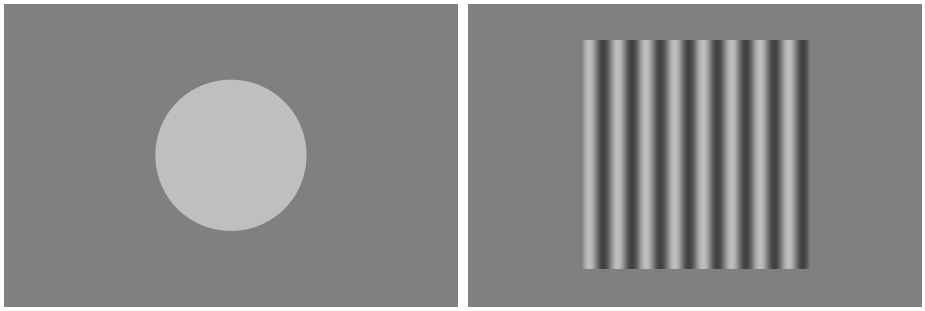

Figure 5.2: Test patterns used in detection threshold experiments. Left: flashing disk, used in Blackwell's and Ferwerda's experiments. Right: static sinusoidal grating used in the definition of the DICOM GSDF.

One perceptually uniform luminance scale, used for the display of radiographic images, is the *DICOM Grayscale Standard Display Function* (GSDF) [52]. The test target used to define the JND is a square with a side length of 2 degrees of visual field, filled with a static horizontal or vertical grating with sinusoidal modulation of 4 cycles per degree and placed on a uniform background with a luminance equal to the mean luminance of the sinusoid. This kind of target is different from the flashing disk used in the detection threshold experiments by Blackwell [14] and Ferwerda [20] described in Section 1.6. The JND is defined as the amplitude which is detected in approximately half the test trials by an average observer. Barten [5, 6, 7] introduced a mathematical model for the contrast sensitivity of the human eye, which is in good agreement with the experimental data and is used in the DICOM standard to compute the TVI curve analytically. The JNDs computed with Barten's model are then accumulated in order to construct a perceptually uniform luminance scale. In order to simplify the computation, a luminance range between 0.05 and 4000 cd/m^2 is selected and the curve in this interval is approximated by an eighth degree polynomial. These limits were chosen in order to contain the typical luminance range of a film-based radiographic image viewed on a light box. The number of JNDs in this interval according to Barten's model is 1023.

5.3 Extension of the DICOM GSDF

As outlined in the introduction, high dynamic range (HDR) displays – such as the ones based on dual-layer LCD technology [76, 23] – can reproduce luminance levels outside the interval on which the DICOM curve is defined, and in particular below its inferior limit. Outside this interval, the polynomial approximation diverges and is therefore inapplicable. It is therefore necessary to investigate how the curve can be extended to low luminance levels.

Apparently, the most intuitive method of extending the DICOM GSDF to a wider luminance range is to use the original curve obtained from Barten's model rather than its polynomial approximation. The algorithm described in the DICOM standard [52] contains a few errors and

inconsistencies; after correcting them, we were able to compute a curve which agrees with the published polynomial approximation for $0.05 \leq L \leq 4000$ cd/m^2 and is consistently defined also outside this interval. Unfortunately, the validity of Barten's model for very low luminance levels is objectionable. In particular, as the luminance L tends to zero, the TVI curve also tends to zero (approximately as $\sqrt{L}$, in agreement with a model known as the DeVries-Rose law), rather than to a finite asymptote like in Blackwell's or Ferwerda's models. The different behavior is clearly visible in the plot of Figure 5.4. Although the comparison might be inappropriate, because the models use different test patterns to measure the JND, we believe that an extension of the DICOM curve based on Barten's model overestimates the visibility of details in dark areas. Indeed, we have verified experimentally that this curve performs poorly when used to decode a medical image for an HDR display, in the sense that a significant portion of the image is rendered with a very low luminance and the visibility of the details is compromised.

Better results are obtained using Blackwell's or Ferwerda's models. If the latter is used, we suggest to use only the photopic TVI curve (1.10), because in a real environment the observer's eye is under-adapted due to the light coming from the ambient and from the bright areas of the image itself, whereas scotopic vision is used only after a long period of adaptation.

5.4 Computational method

We propose [30] a simple algorithm which generates a perceptually uniform scale of n luminance levels between two arbitrary extremal values $L_{\min}$ and $L_{\max}$. More precisely, the algorithm outputs an array of n monotonically increasing luminance levels $L_1, L_2, \ldots, L_n$, such that $L_1 = L_{\min}$, $L_n = L_{\max}$, and the luminance difference between two adjacent levels L_j and L_{j+1} is a constant multiple of the JND as defined in Equation (5.3). In a typical application, the extremal values $L_{\min}$ and $L_{\max}$ correspond to the black and white levels of the display and the number of levels n is determined by the bit depth of the device. The algorithm we propose operates as follows:

1. As a first starting guess, generate n logarithmically-spaced luminance levels between $L_{\min}$ and $L_{\max}$.

2. Compute the perceptual distance $\Delta P(L_j, L_{j+1})$ between each pair of adjacent levels using the expression (5.1).

3. Compute the perceptual brightness P_j corresponding to each level L_j by accumulating the perceptual distances ΔP of the previous step using the expression (5.2). The perceptual brightness values P_j obtained at this stage are unevenly spaced in general.

4. In order to build a perceptually uniform scale, generate n uniformly spaced perceptual brightness values $\hat{P}_j$ between 0 and the extremal value P_n found at the previous step and compute the corresponding luminance levels $\hat{L}_j$ by performing an inverse piecewise-linear interpolation.

A graphical representation of the operations performed by the algorithm is reproduced in Figure 5.3. The luminance levels found in this way have, to a good approximation, an equal perceptual distance, besides a small error deriving from the interpolation. If extra precision is required,

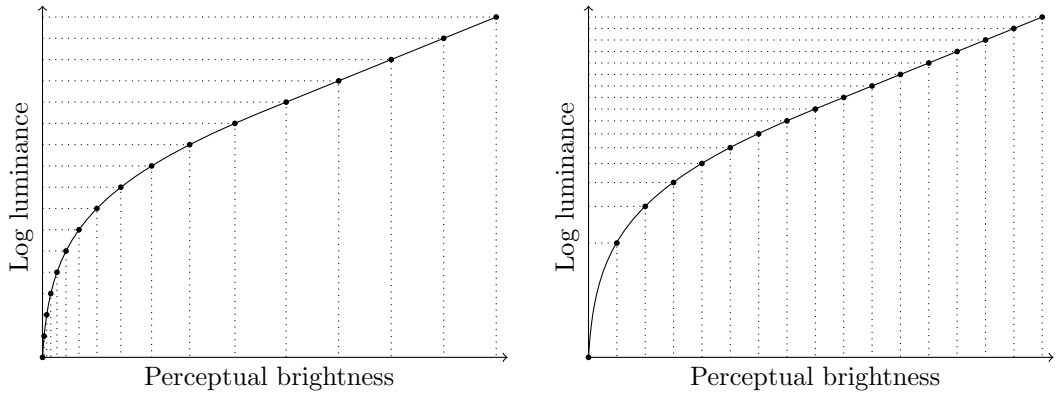

Figure 5.3: Graphical representation of the operations performed by the proposed algorithm. Left: step 2 (luminance to perceptual brightness). Right: step 3 (perceptual brightness to luminance).

it is possible to iterate the steps 2 – 4 in the algorithm. Alternatively, it is possible to refine the solution with a simple fixed point iteration. By adding all the equations in (5.3), for $j = 1 \dots n-1$, we obtain the expression

$$L_n - L_1 = k \left[\frac{1}{2} \operatorname{TVI}(L_1) + \sum_{j=2}^{n-1} \operatorname{TVI}(L_j) + \frac{1}{2} \operatorname{TVI}(L_n) \right], \tag{5.4}$$

which is solved for k. The luminance levels are then updated by solving the equations (5.3) for the L_j, using this new value of k and the TVI values computed in the previous iteration. It is easy to see that this operation consists in the computation of a simple cumulative sum. We have verified that a single iteration is enough to produce a satisfactory precision, which is not a strict requirement anyway because the JND is a statistical quantity as previously noticed and the TVI function is an analytical approximation of experimental data.

5.5 Ambient light and under-adaptation correction

Besides the extension of the mapping curves, one issue encountered when using high dynamic range displays is that the luminance emitted by the device in the dark portions of the images can be lower than the ambient light L_{amb} reflected by the screen. A simple modification of the proposed algorithm is able to generate a curve which compensates this effect by amplifying appropriately the details in dark areas.

In our model we assume that the luminance which reaches the eye is the sum of the luminance emitted by the display and a constant offset L_{amb}. We seek a scale of *display* luminances which will have an equal perceptual distance after the offset is added. By modifying Equation (5.3), we obtain:

$$(L_{j+1} + \cancel{L_{\text{amb}}}) - (L_j + \cancel{L_{\text{amb}}}) = \frac{k}{2} \left[\operatorname{TVI}(L_j + L_{\text{amb}}) + \operatorname{TVI}(L_{j+1} + L_{\text{amb}}) \right] \tag{5.5}$$

It is easy to see that the corrected levels can be computed by replacing $\operatorname{TVI}(L)$ with $\operatorname{TVI}(L + L_{\text{amb}})$ in the algorithm described in Section 5.4.

This modified method was tested by decoding and displaying – on an HDR display based on a dual layer LCD technology [23] – synthetic test images containing details with equal perceptual amplitudes over different background levels. Ideally, all the details should be equally visible and appear with the same intensity, regardless the background. This was not the case with the extended DICOM curve based on Barten's model, and indeed the details in the dark portions of the image were less visible or invisible at all. Their visibility was improved when the mapping curves based on Blackwell's or Ferwerda's models were used, and further improved when the ambient light compensation method was turned on.

One source of error is also due to the fact that the light from the bright areas of the image is scattered, both by the veiling glare inside the observer's eye and by optical crosstalk mechanisms inside the display; this scattered light further reduces the visibility of the details in the dark portions. Indeed, we have verified experimentally that the value of L_{amb} which produces a visually pleasing image is significantly higher than the actual ambient light reflected by the screen, whose value was obtained by measuring with a radiometer the luminance of the screen when the display is turned off. In order to compensate this *under-adaptation* effect, the value of L_{amb} should be made space-varying and image-dependent in general. We are currently investigating methods which attempt to estimate the veiling glare component in an image and to adapt the mapping curve locally.

5.6 Results

A plot of the TVI functions computed using Barten's, Blackwell's and Ferwerda's models can be seen in Figure 5.4. It is immediately clear that Barten's curve is significantly different from the other two, and in particular yields lower values for the JND, especially in dark areas. The resulting mapping curves, computed with the algorithm described in Section 5.4 with $L_{\text{min}} = 0.01$ cd/m^2 and $L_{\text{max}} = 1000$ cd/m^2, are compared in Figure 5.5. The plot shows that the curve based on Barten's model (which corresponds to the DICOM GSDF used in medical imaging) produces the darkest image as previously explained.

The effect of the ambient light compensation can be seen in the plot of Figure 5.6. We used Barten's model with the same parameters used in the previous plots. The ambient light compensation has the effect of increasing the curve slope for low luminance levels, and consequently amplifying the details in dark areas. It should be noticed that, since the values of L_{min} and L_{max} are fixed, an amplification of the details in the dark areas inevitably introduces a compression in the bright areas. Therefore, the improvement in visibility comes at a cost and a suitable trade-off should be found.

As previously noticed, the results can only be appreciated by viewing the mapped images on a calibrated high dynamic range display and can not be reproduced on a conventional display or on print. For this reason, we did not include sample images in this chapter.

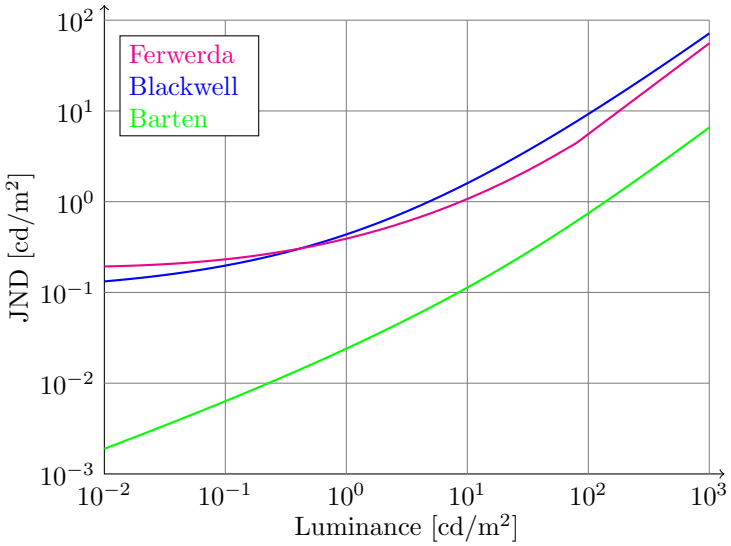

Figure 5.4: Comparison of the TVI functions obtained using Barten's, Blackwell's and Ferwerda's models

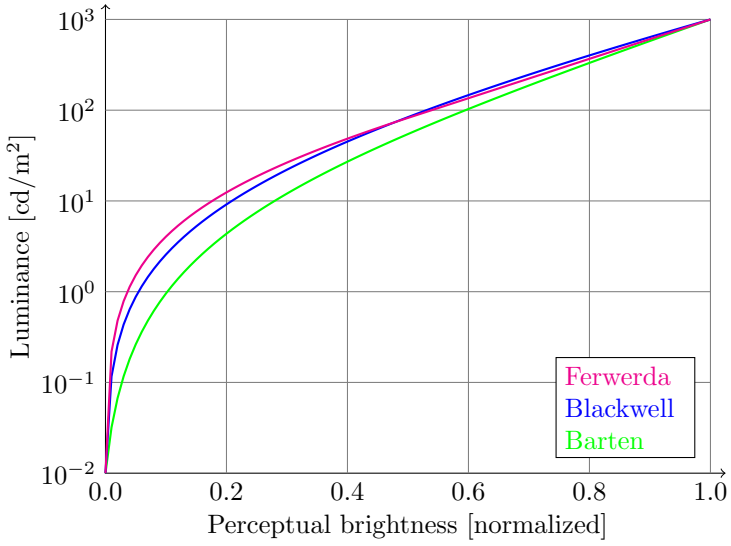

Figure 5.5: Comparison of the mapping curves generated by the proposed algorithm using Barten's, Blackwell's and Ferwerda's TVI functions.

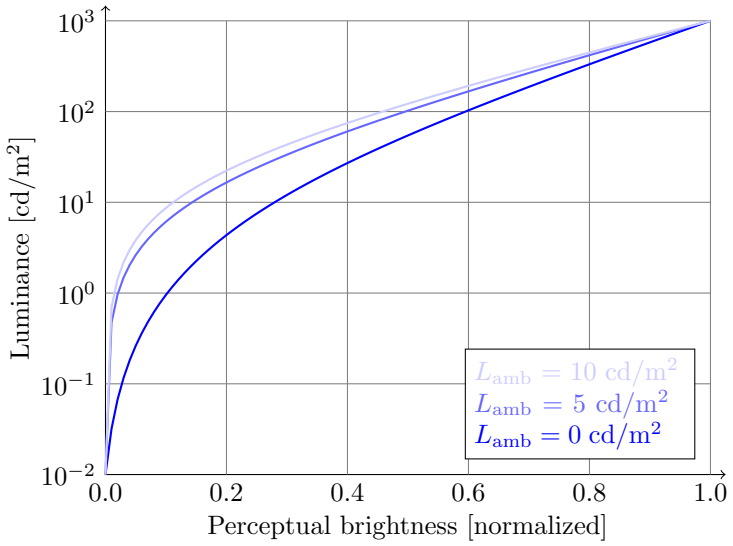

Figure 5.6: Comparison of the mapping functions obtained with different values of ambient light compensation.

Chapter 6

Dual layer LCD display for medical applications

6.1 Introduction and motivation

Until recent times, a diagnosis in radiology typically involved the examination of the x-ray film on the light box. An image viewed in this way can reach a peak brightness of around 4000 cd/m^2 and an overall luminance range of 3 to 4 orders of magnitude. This performance allows the radiologists to detect the presence, in the image, of low-density details which introduce a very small luminance difference but carry relevant clinical information. The theory underlying the definition of the DICOM grayscale standard display function, described in Chapter 5, shows that an average observer is able to discriminate approximately 1023 visible levels in the range between 0.05 and 4000 cd/m^2. The transition from conventional film to digital modalities can introduce several benefits, including reduced diagnosis times due to the near-instant visualization of the image, ease of archival and retrieval, and possibility of using a wide range of image processing and enhancement techniques. Current digital x-ray detectors are already able to record data with a depth of 14 to 16 bits. Unfortunately, the display technology still represents a bottleneck because most current devices are only able to reproduce a significantly lower luminance range than the conventional radiographic film. Cathode ray tube (CRT) displays offer a very low black level, but their resolution and peak brightness are limited by a physical issue: if the intensity of the cathode ray is increased in order to produce a brighter image, the beam size also increases and the resulting image becomes blurred. Moreover, CRTs can introduce geometric distortions in the image, their performance tends to degrade with age and the image flicker can introduce eye fatigue unless a very high refresh rate is used. Liquid crystal displays (LCDs), on the other hand, offer a high resolution with no geometric distortion, and their peak brightness only depends on the power of the backlight unit, being limited only by the efficiency of the heat dissipation. Advantages of LCDs also include a flat surface and small physical dimensions, no flicker, good longevity, low power consumption and low electromagnetic emissions. However, LCD displays have a poor dark level performance due to technological limitations. Conventional LCD displays are constituted of two basic elements: a uniform light source – usually a Cold Cathode Fluorescent Lamp (CCFL) backlight unit – and the actual liquid crystal panel. Depending on the applied voltage, the light transmittance of each cell in the panel is modulated from a maximum value $T_{\max}$ (about 0.3 for monochrome LCD panels and 0.1 for color panels) to a theoretically off-state value $T_{\min}$. The *dynamic*

range or *static contrast ratio*, defined as the ratio between the white level and the black level measured simultaneously on a single frame, is equal to $T_{\max}/T_{\min}$ and is independent of the backlight. When a cell is in its off state, some amount of light is still able to leak through: with current technology, the static contrast of a medical-grade LCD panel is limited to a range of $500 \div 800$. As previously mentioned, this could represent a bottleneck in specific applications such as mammography, where the diagnosis is linked to the detection of very small details with fine luminance differences. Moreover, the driving electronics of most LCD panels allows a bit depth of 8 bits, i.e. the visualization of 256 distinct luminance levels; for medical applications, some manufacturers have presented 10 bit panel prototypes. It should be noticed that the bit depth of an LCD panel does not determine the dynamic range, which depends on the values of $T_{\max}$ and $T_{\min}$, but rather the number of discrete gray levels that can be reproduced within this interval; in general, the transmittance of a liquid crystal cell is a nonlinear function of the applied digital driving level.

This Chapter is structured as follows. In Section 6.2 we introduce a prototype high dynamic range display based on the dual layer LCD prototype; we describe the hardware construction and mode of operation and compare this solution with an alternative design which has been recently proposed. In Section 6.3 we describe the luminance properties of the prototype and illustrate the techniques we used to perform the characterization. In Section 6.4 we introduce the concept of splitting algorithms justifying their need and illustrating their objectives. In the Sections 6.5 – 6.9 we describe in detail the different splitting algorithms we propose, which add incrementally additional and more complex requirements. In Section 6.10 we present some experimental results obtained using the proposed method and discuss the influence of the algorithm parameters. In Section 6.11 we illustrate the phenomenon of in-panel glare caused by an unwanted scattering of light between the panels; despite the very small magnitude of the consequent artifacts, we propose a modification in the splitting algorithm which addresses this issue. Finally, in Section 6.12 we present the conclusions and suggest possible future work on this subject.

6.2 Dual layer LCD prototypes

Dynamic range reduction algorithms, such as the ones described in Chapter 3, are able to reduce the dynamic range of an image without compromising the visibility of fine details. However, their use in the medical field is discouraged because the photometric distortion that is intrinsically introduced by the processing may cause an incorrect diagnosis. In particular, physicians often request that points corresponding to the same tissue density be reproduced with the same luminance, and only a global tone mapping curve is able to guarantee this correspondence.

Thanks to recent technological advancements, high dynamic range LCD displays are beginning to appear on the market. The black level in an HDR-LCD display is typically reduced by modulating the light coming from the backlight unit a second time, in order to dim it locally in correspondence of the dark areas of the image. A recently proposed architecture [67, 74] uses an array of individually modulated LEDs as a backlight unit. The LED backlight has a significantly lower resolution than the LCD panel, and this creates the need of an appropriate processing of the image displayed on the latter. Moreover, this solution is costly with regard to the high power consumption and consequent thermal problems, together with the complex

control of the backlight system. One alternative solution is to stack two LCD panels one on top of the other over a uniform backlight unit, thus obtaining a *Dual Layer* (DL) LCD [76]. In this way, the global transmittance is the pointwise product of the transmittances of the two individual layers and the resulting dynamic range is theoretically squared for two equal LCD panels. The grayscale reproduction accuracy is also increased since the light coming from the backlight unit is modulated twice. It should be precised, however, that the resulting bit depth is not doubled because, due to the commutative law of multiplication, different combinations of the two panels can produce the same output level; the actual concept of bit depth can not be defined unambiguously because the luminance levels obtained by the combination of the two layers are partially overlapped and unevenly spaced, and this issue is examined in detail in 7. As a drawback, the total transmittance of the DL-LCD is reduced, therefore a higher-brightness backlight unit must be used to obtain the same white luminance. Furthermore, the planes on which the two images are formed have a finite distance due to the presence of the supporting glass and polarizer filters in between, and this introduces a *parallax* effect when the display is viewed off-axis, even with perfectly aligned LCD panels. In order to reduce the visibility of the parallax error, an appropriate *splitting algorithm* is needed to generate the two images which drive the panels.

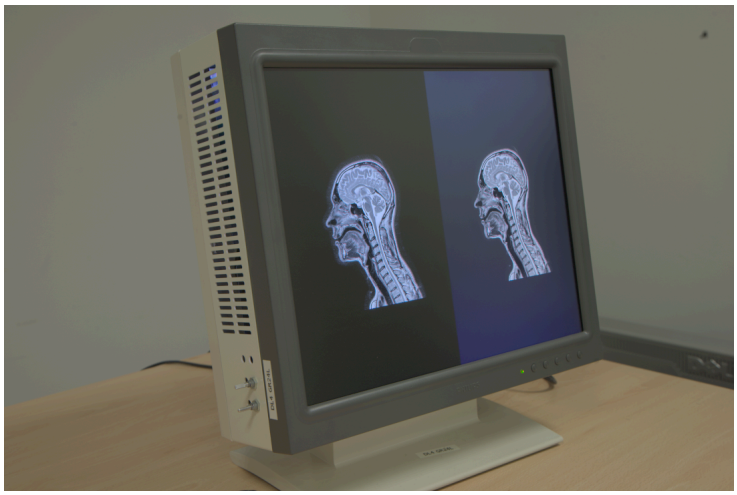

Figure 6.1: Photograph of the prototype Dual Layer LCD display developed by FIMI-Philips. On the left half of the screen, a full dual layer image is displayed. On the right half, the appearance of a conventional LCD display is simulated by displaying the image on the frontpanel only over a uniform white backpanel. Due to the high dynamic range of the prototype, the original photograph has been tone mapped for display and print.

Several prototypes of dual layer LCD displays for medical imaging applications were built at FIMI-Philips and are currently under test. Each prototype was assembled using two identical and commercially available 18-inch grayscale LCD panels with a resolution of 1280×1024 pixels. The panels employed are from CMO / IDTech and use the Super-In-Plane Switching pixel design (S-IPS) for good viewing angle characteristics. After removing the original backlight unit, the individual LCD glass panels are mechanically matched one on top of the other, taking care to align the active areas at best. The panel facing the backlight unit is referred to as the *backpanel*, and the one seen by the viewer is referred to as the *frontpanel*. Since the two polarizers of each panel are oriented in orthogonal directions, the back panel is flipped around a vertical axis

as shown in Figure 6.2. In order to minimize the parallax effect, the two panels are placed together as close as possible with a special mechanical construction: with the IDTech panels used in the prototypes, the distance between the liquid crystal cells is approximately 1.9 mm; one of the two inner polarizers is redundant and could be removed, thus reducing the distance to 1.6 mm. In a future product, a dedicated dual layer LCD unit could be designed in order to reduce this distance to the minimum value. Finally, the two panels are combined to a single high-brightness backlight unit made of 12 CCFL lamps directly illuminating the LCD. The backlight used in the current prototype is calibrated for a maximum luminance of 500 cd/m^2 and stabilized via feedback from a photodiode. The driving electronics is properly arranged in order to allow each panel to be driven independently from a standard DVI input. To operate the display a standard PC is used, on which a graphics card with dual DVI output is installed, and each panel of the dual layer monitor is attached to one output.

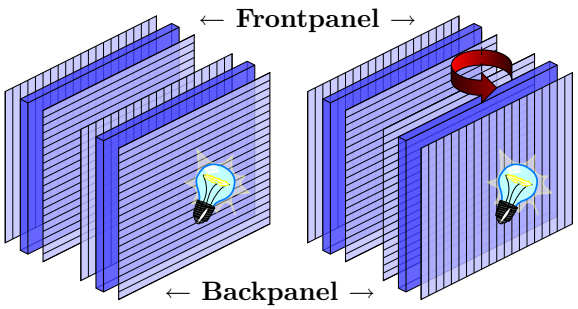

Figure 6.2: Alignment of the displays. On the left, both displays are stacked with equal orientation. As a result, the two inner polarizers are oriented perpendicularly with respect to each other, such that all light will be blocked. In the right figure, the backpanel has been rotated around its vertical axis such that the inner polarizers are now aligned.

A picture of a prototype is visible in Figure 6.1. On the left half of the screen, the dual layer display is fully exploited; for comparison, on the right half of the screen a conventional single layer LCD is simulated by displaying the whole image on the frontpanel and a white background on the backpanel. The improved black level of the dual layer display is clearly visible: its value, measured in a dark environment, is below 0.01 cd/m^2, corresponding to a dynamic range of over 50000. Besides the high dynamic range and the improved bit depth, the grayscale dual layer LCD prototypes also features an enhanced performance in terms of viewing angle, in the sense that a reduced luminance drop-off with angle compared to a conventional display was measured. In Figure 6.3 the prototype is photographed next to a standard monitor in order to show this improvement.

Color prototypes have also been built by stacking a color frontpanel over a grayscale backpanel. Color LCD panels, however, have a lower transmittance than the grayscale ones due to the presence of the color filters; this causes the need of a higher-brightness backlight unit which can be critical with regard to heat dissipation. The use of a color backpanel is not advisable, since it would further reduce the peak brightness and likely introduce moiré artifacts without introducing a significant benefit. Besides having a lower peak brightness, color prototypes also have a smaller importance in the current medical applications, since a high dynamic range is typically required in radiology, which mostly involves the display of grayscale data.

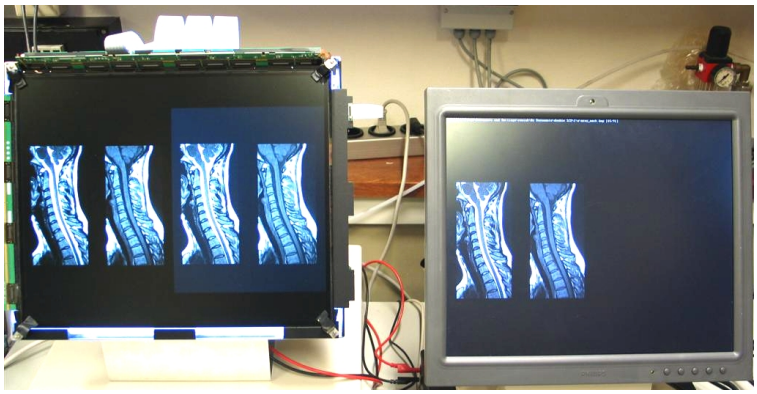

Figure 6.3: Picture of the prototype. Clearly visible is the very good black level in the left part of the left display (the dual layer LCD prototype). The right side of the prototype screen shows the image as it would look on a normal greyscale LCD monitor, which is placed to the right of the prototype for comparison. The right photo shows that also the viewing angle is greatly improved.

6.3 Luminance characterization

As previously mentioned, the luminance output $L_{\text{out}}(x, y)$ of a dual layer LCD at pixel location (x, y) is equal to

$$L_{\text{out}}(x, y) = B\, T_{\text{b}}(x, y)\, T_{\text{f}}(x, y), \tag{6.1}$$

where B is the backlight intensity, measured in cd/m^2, and $T_{\text{b}}(x, t)$ and $T_{\text{f}}(x, y)$ are the transmittances of the two panels, which are adimensional. This multiplicative law, which follows intuitively from the structure of the display, was also proved by experimental measurements.

In order to characterize the luminance output of the dual layer prototype as a function of the digital driving levels (DDLs) of the two panels, we designed a software application which displays a sequence of uniform full-screen patches on the two panels and reads the corresponding luminance output with a luminance meter. A first set of measurements was performed in the laboratories of FIMI-Philips using a Konica Minolta CS-200 colorimeter. The instrument is connected with an USB cable to the PC which drives the display and runs the measurement program, and a vendor-supplied dynamic link library allows our code to read and process the data. The DDLs of the patch sequence are read from a text file; after each patch is displayed, the program waits for a specified time delay in order to allow the luminance to settle and then starts the measurement. The colorimeter adjusts the integration time automatically and transmits the measured data as a formatted text string. In order to compensate possible variations of the backlight intensity, caused for instance by a temperature drift, we displayed and measured periodically a white patch. The dynamic range of the CS-200 colorimeter is not sufficient to measure the darkest levels of the dual layer display, therefore we could only measure the look-up-tables (LUTs) of the individual panels. The measurements showed that different specimens can have significantly different LUTs, therefore each prototype should be characterized individually in order to obtain the most accurate image rendition. Moreover, the darkest levels exhibit a chrominance variation and become bluish; this effect is a characteristic of the in-plane-switching technology and is also visible by naked eye.

A more detailed set of measurements was performed in the laboratories of the Center for

Devices and Radiological Health, US Food and Drug Administration, using an International Light IL1700 radiometer capable of measuring light intensities in a range of over 10 orders of magnitude. The radiometer samples the current generated by the photodiode in the probe every 0.5 seconds and transmits an ASCII-encoded text string through an RS-232 serial cable. We used a modified version of our measurement program which reads the data coming from the serial port and takes the average of multiple samples in order to reduce the measurement noise; the number of samples is adjusted automatically as a function of the mean and variance of the measurements in order to obtain a signal-to-noise ratio above a specified threshold. A measurement of the entire dual-layer LUT, with all the 65536 possible combinations of frontpanel and backpanel DDLs, would take several days and is impractical to perform even on a single specimen. Therefore, we made a partial measurement by displaying all the 256 levels on the frontpanel and a constant level on the backpanel. The measurement was repeated with 6 different values for the backpanel; the process was then repeated swapping the two panel signals. The luminance corresponding to the DDL pairs which were not measured was estimated with an appropriate interpolation. If the multiplicative law of Equation (6.1) holds exactly, the logarithm of the measured luminance as a function of the backpanel and frontpanel DDLs can be separated into a sum of two functions each depending on one variable only. In practice this identity can not hold exactly due to measurement errors, especially for the darkest luminance levels where the dark current of the photodiode, the aforementioned chrominance variation and even small amounts of ambient light can influence the measurement. We seek therefore a least-squares approximation, and a function which satisfies this property for a given boundary is the Coons surface [65]. A plot of the dual layer LUT estimated with this method is shown in Figure 6.4.

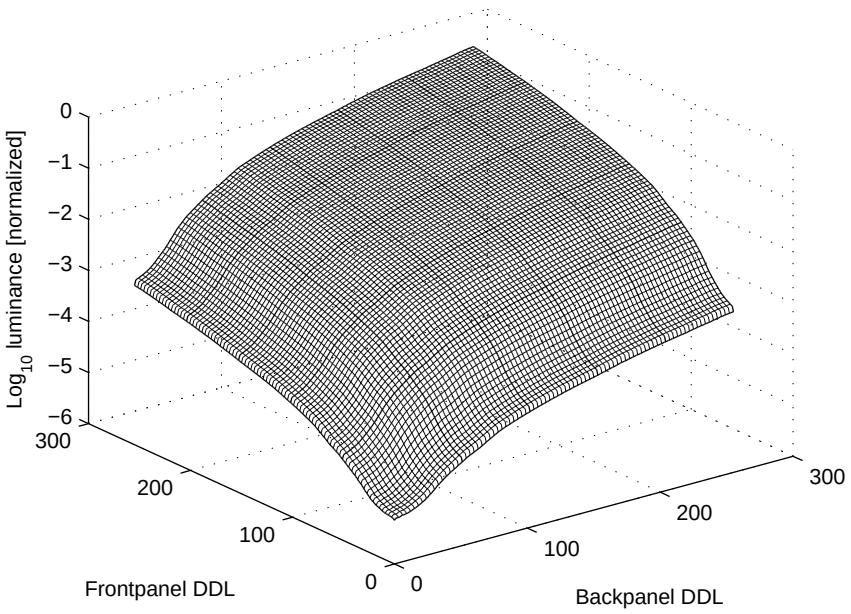

Figure 6.4: Output luminance of a dual layer display as a function of the backpanel and frontpanel DDLs. The thick lines represent the measured values; the thin lines represent the interpolated values.

6.4 Splitting algorithms

The processing used to generate the two images to be displayed on the two panels plays a fundamental role in the performance of the device [23]. The simplest possible technique is to perform the splitting on a pixel-by-pixel basis. More precisely, if we indicate with $L_{\text{in}}(x, y)$ the luminance of the input image at pixel location (x, y) and follow the notation used in Equation (6.1), the splitting algorithm takes the form

$$T_{\text{b}}(x, y) = F\big(L_{\text{in}}(x, y)\big) \qquad T_{\text{f}}(x, y) = \frac{L_{\text{in}}(x, y)}{B\, T_{\text{b}}(x, y)}. \tag{6.2}$$

In other words, the backpanel is computed by mapping the input luminance with a suitable nonlinear function $F(\cdot)$, and the frontpanel is subsequently computed by division in order to guarantee that the product of the two images reproduces the input. The splitting is computed on linear data; the nonlinear encoding of the source image (if present) and the response of the liquid crystal panels are compensated appropriately by mapping the data before and after the processing. An intuitive choice for the function $F(\cdot)$, suggested by the symmetry of the system, is a square root [58]: in this way, each panel displays the same image. However, in a real device the liquid crystals are enclosed between two glass plates which introduce a finite distance between the planes on which the images are formed even if the two panels are in contact. Instead of seeing the correct image $L_{\text{out}}(x, y) \triangleq B\, T_{\text{b}}(x, y)\, T_{\text{f}}(x, y)$, an observer looking at the display from an off-axis position sees a distorted image $\tilde{L}_{\text{out}}(x, y) \triangleq B\, T_{\text{b}}(x + \Delta x, y + \Delta y)\, T_{\text{f}}(x, y)$, where the displacements Δx and Δy depend on his viewing angle. This form of distortion is referred to as *parallax error* and gives rise to artifacts such as those shown in Figure 6.5. It is convenient to measure the distortion in the perceived image with the relative error:

$$E(x, y) \triangleq \frac{\tilde{L}_{\text{out}}(x, y) - L_{\text{out}}(x, y)}{L_{\text{out}}(x, y)} = \frac{T_{\text{b}}(x + \Delta x, y + \Delta y) - T_{\text{b}}(x, y)}{T_{\text{b}}(x, y)}. \tag{6.3}$$

With this choice, an expression is derived which depends on the backpanel only. The use of the relative error also has a perceptual motivation since the perception of luminance by the human visual system approximately follows Weber's law. Weber's law only holds in a limited range of luminance levels; however, a more accurate measure would yield an expression which also depends on the frontpanel and greatly increase the problem complexity.

In order to reduce the distortion, the backpanel image should be blurred, so that a small displacement does not alter the pixel values excessively. The frontpanel image is then sharpened to compensate for the blurring. Designing the corresponding algorithm is however a nontrivial task which should meet different and conflicting requirements:

- **Feasibility**: The transmittances computed by the splitting algorithm should be greater than the panel black level and smaller than the panel white level.
- **Perfect reconstruction**: The image resulting from the combination of the backpanel and the frontpanel should be equal to the input image when the two panels are aligned. If this constraint can not be satisfied, the splitting algorithm should be designed in order to minimize the visibility of the reconstruction error.
- **Parallax error reduction**: The splitting algorithm should exploit the available degrees of freedom in order to minimize the reconstruction error when the two panels are viewed from a misaligned direction.

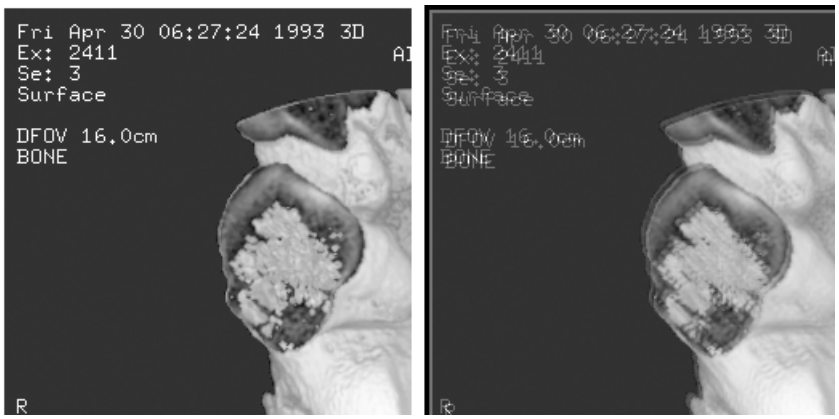

Figure 6.5: Example of parallax error with square root splitting. Left: original image. Right: simulation of the displayed image with a horizontal and vertical shift of 3 pixels in the backpanel.

- **Computational efficiency**: The computational complexity of the splitting algorithm should be low enough to allow real-time processing on a standard PC or on low-cost graphics hardware.

In the first stage of our research [23], we considered a class of splitting algorithms which meet the perfect reconstruction constraint:

$$B\,T_{\mathrm{b}}(x,y)\,T_{\mathrm{f}}(x,y) = L_{\mathrm{in}}(x,y) \qquad \forall (x,y). \tag{6.4}$$

This simplifies the problem because only the backpanel needs to be computed; the frontpanel is then generated automatically by division. It is also possible to compute the frontpanel first, but the former approach is more natural to follow because the design objectives are easier to express for the backpanel. After the splitting, the frontpanel and backpanel luminances are converted to digital driving levels (DDLs) and quantized to 8 bits. As previously noticed, a nonlinear mapping must be performed in order to compensate the distortion introduced by the panels; this operation is commonly known as *gamma correction*. For greater accuracy, we measured the actual response of the panels used in the prototypes, rather than using an analytical curve defined in standard recommendations for display or television devices [37]. The pixels of grayscale LCD panels are actually made of three subpixels which can be independently driven by the three RGB channels of the DVI input. This permits the use of *subpixel dithering* in order to improve the perceived bit depth of the device.

A splitting algorithm can be extended to color images with only minor modifications. The splitting is computed on the luminance only, since a grayscale backpanel must be generated in any case; the frontpanel is then computed by dividing each of the RGB channels of the input image by the backpanel. It is convenient to compute the luminance by taking the maximum of the three RGB channels, following the definition used in the HSV color space. With this choice, the three channels of the frontpanel image are upper bounded by the corresponding grayscale image and it is easier to prevent unwanted clipping.

A block diagram of a generic perfect reconstruction algorithm is shown in Figure 6.6. The critical part in the algorithm is clearly the computation of the backpanel, and in the following we will describe some possible solutions.

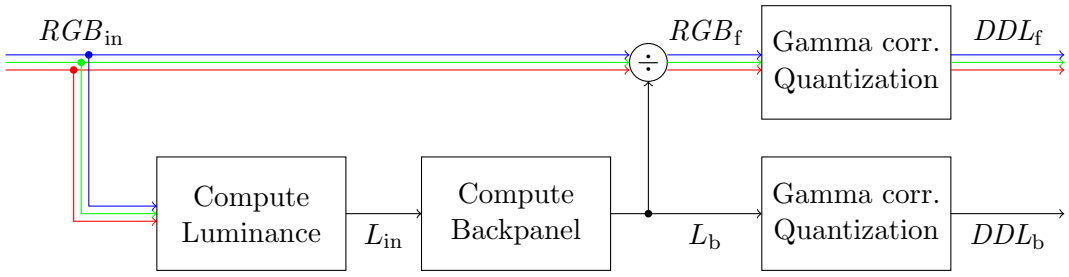

Figure 6.6: Block diagram of a perfect-reconstruction splitting algorithm.

6.5 Method 1: Constrained filtering of the backpanel

A simple improvement on the square root technique described in the previous section consists in blurring the square root of the input image with a lowpass filter in order to obtain a smooth backpanel. The frontpanel is then computed by division, following the scheme of Figure 6.6:

$$T_{\text{b}}(x,y) = \text{BLUR}\left(\sqrt{c\,L_{\text{in}}(x,y)}\right) \qquad T_{\text{f}}(x,y) = \frac{L_{\text{in}}(x,y)}{B\,T_{\text{b}}(x,y)}, \tag{6.5}$$

where c is a suitable scaling factor, with the physical dimension of m^2/cd, used to convert from luminance to transmittance. Blurring is however a nontrivial task, because it must take into account the limited dynamic range of the panels. If a linear lowpass filter is used, a portion of the pixels in the backpanel will be darker than $\sqrt{c\,L_{\text{in}}(x,y)}$ and the corresponding pixels in the frontpanel will be brighter; if no precaution is taken, some pixels in the frontpanel might exceed the white level. An example of this problem on a 1D test signal is shown in Figure 6.7. In order to prevent the clipping of the frontpanel, some sort of nonlinear filter must be used.

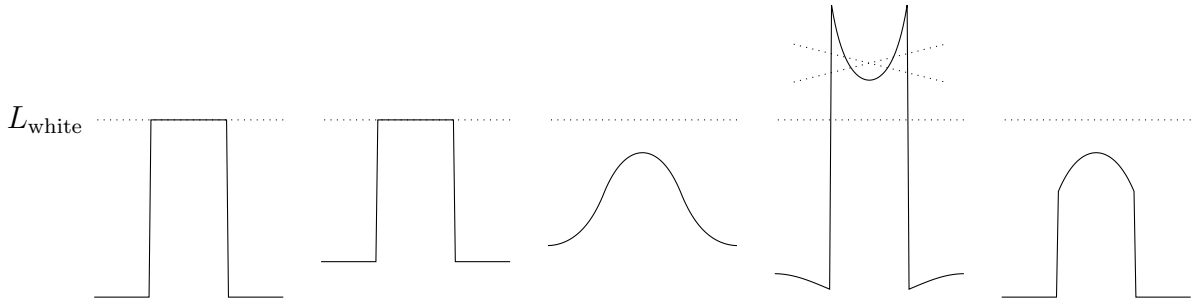

Figure 6.7: Saturation artifacts. Left to right: input signal, square root, backpanel, frontpanel with clipping, reconstructed signal.

A problem which often arises in signal and image processing is the computation of *constrained lowpass filters*, or *envelopes*. This operation basically consists in approximating a signal with a smoother function which is bounded from below by the signal itself. The effect can be visually interpreted as an elastic rope or membrane resting on top of the signal graph. If a higher smoothness is required, it is possible to simulate the behavior of an elastic beam or plate, typically by means of spline models. A comparison of the two formulations on a 1D signal is illustrated in Figure 6.8; the main difference is that a membrane-based envelope can have a discontinuous first derivative in the points where the constraint is active, whereas spline-based envelopes are smooth also in this case.

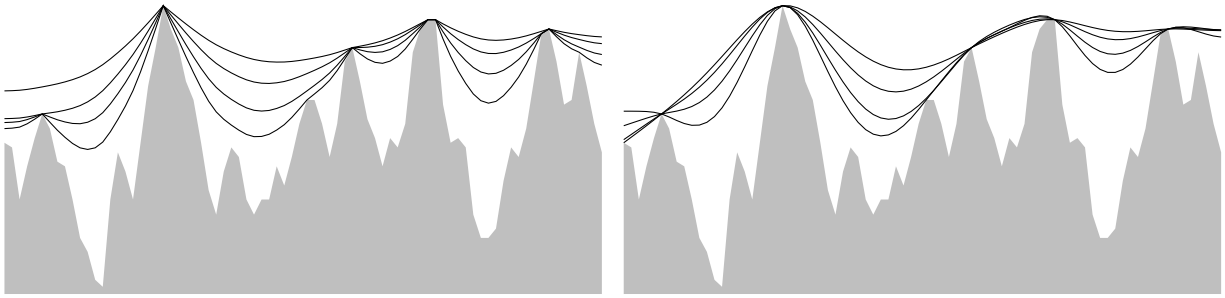

Figure 6.8: Membrane-based and spline-based envelope filters with different "bandwidths".

A possible technique which avoids the clipping of the frontpanel consists in filtering the square root of the input image with a constrained lowpass filter rather than a simple linear lowpass filter. The backpanel computed in this way meets by definition the constraint $T_\mathrm{b}(x,y) \geq \sqrt{c\,L_\mathrm{in}(x,y)}$. Consequently, the frontpanel computed by division will be upper bounded by the square root, which is lower than the white level provided that the input image is suitably scaled, and clipping is avoided:

$$T_\mathrm{b}(x,y) \geq \sqrt{c\,L_\mathrm{in}(x,y)} \qquad \Rightarrow \qquad T_\mathrm{f}(x,y) = \frac{L_\mathrm{in}(x,y)}{B\,T_\mathrm{b}(x,y)} \leq \frac{\sqrt{L_\mathrm{in}(x,y)}}{B\sqrt{c}} \tag{6.6}$$

This property is visualized in Figure 6.9 using a simple 1D test signal. An application to a real image is shown in Figure 6.10.

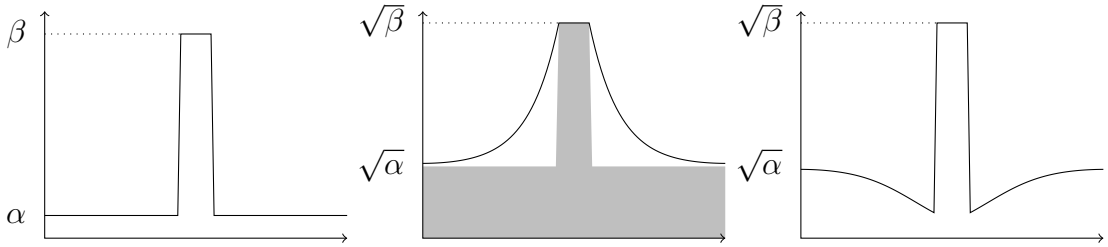

Figure 6.9: Example of splitting for a 1D test signal. Left to right: input, backpanel, frontpanel.

A constrained lowpass filter can be implemented in several ways. A heuristic approach may consist in blurring the image with a linear lowpass filter and then adding an offset to the output or parts of it. More advanced methods may formulate the filtering operation as a constrained optimization problem. A possible objective function is

$$\iint \Big\{ \big\|\nabla f(x,y)\big\|^2 + \lambda \big[f(x,y) - u(x,y)\big]^2 \Big\}\, dx\, dy = \mathrm{Min} \qquad f(x,y) \geq u(x,y), \tag{6.7}$$

where $u(x,y)$ is the input image, $f(x,y)$ is the filtered output and the integral is computed over the whole image area. The first term privileges a smooth output, and the second term privileges an output which closely approximates the input; the scalar λ allows to set the tradeoff between these two objectives and adjust the "bandwidth" of the filter. A spline-based filter can be obtained by replacing the first term in (6.7) with a thin plate spline energy functional;

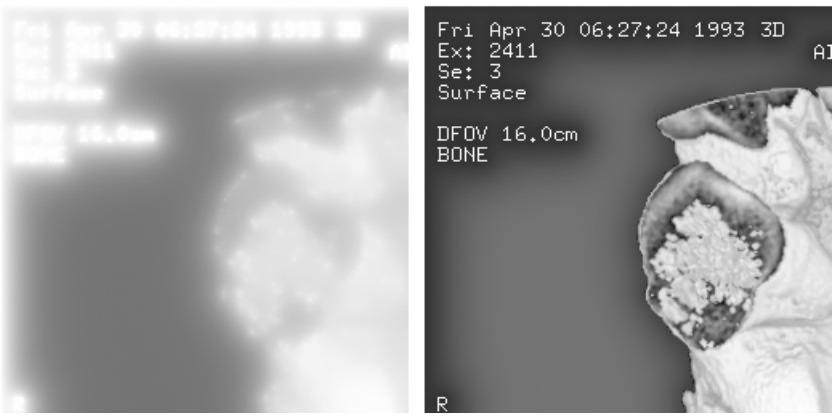

Figure 6.10: Backpanel and frontpanel computed with square root and constrained lowpass filter.

this approach involves higher-order derivatives and has a remarkably higher computational complexity. Once (6.7) is discretized, a quadratic programming problem is obtained, which can be solved by means of appropriate iterative methods that exploit its structure and sparsity.

The computation of constrained lowpass filters typically arises in a class of algorithms for the dynamic range reduction of images based on the Retinex model of the human visual system (HVS) [44, 43]. According to the model, the luminance in an image is given by the pointwise product of the *illumination*, which depends on the light sources, and the *reflectance*, which instead depends on the objects. Studies on the HVS have revealed that the visual sensation mostly depends on the object reflectance, and remains unaltered if the light source illuminating the scene is changed within certain limits. Retinex-based algorithms attempt to decompose the input image into the product of illumination and reflectance; the illumination is then mapped in order to reduce its dynamic range, while the reflectance is left unaltered or optionally enhanced. The illumination is typically a smooth function, and is always larger than the image luminance because the objects reflect only a fraction of the incident light; based on these assumptions, a constrained lowpass filter appears a suitable operator for the estimation of this component. The problem of illumination estimation in Retinex-based algorithms bears some remarkable similarities with the image splitting in a dual layer display, therefore we made some attempts to adapt techniques developed for the former problem to the present application. The algorithm by Frankle and McCann [21], which we described in Section 3.4, can be applied successfully to the present case. Recursive filters have also been investigated [68], but the presence of the constraint can give rise to banding artifacts in the filtered image. Variational methods have been recently proposed [40, 41]. Despite the similar objective, algorithms designed for Retinex applications often have some quality issues which make them inapplicable for the computation of the backpanel in a dual layer display. This happens because, in the former application, artifacts introduced by the constrained lowpass filter generally produce a hardly noticeable error in the output image; therefore, the quality requirements are much lower than in the present case. In particular, many algorithms are unable to completely smooth out sharp edges. However, the illumination component in an image can contain sharp edges in the points where an object covers the light source; therefore, the preservation of edges is a desired feature rather than an artifact, and indeed some formulations of the Retinex algorithm use *edge-preserving* lowpass filters which take this property into account and avoid the formation of unwanted halos.

6.6 Method 2: Generation of the backpanel image by constrained optimization

In the algorithm described above, the constraints have an asymmetric nature, since they only include a lower bound for the backpanel. In presence of dark details on a light background, the constrained lowpass filter tends to "fill the hole" producing a uniform backpanel and the detail is reproduced on the frontpanel only. In case of light details on a dark background, instead, the backpanel must follow the transition as shown previously in Figure 6.9, and on some critical images this variation in the backpanel can give rise to visible parallax error even if a narrowband lowpass filter is used. Artifacts can appear especially around synthetic parts of the image such as text or frames over a uniform dark background. An example is shown in Figure 6.11. Improved implementations of the constrained lowpass filter can reduce this problem, but not avoid it completely, since it is intrinsically present in the method.

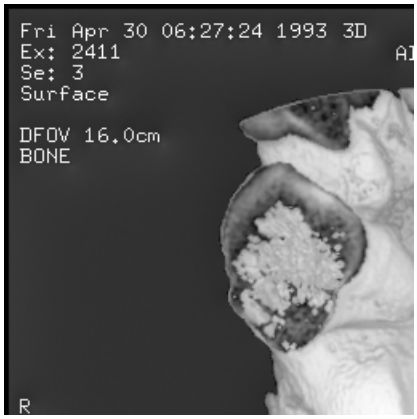

Figure 6.11: Example of parallax error with synthetic images.

Moreover, the constraints themselves can be chosen in a more clever way. By using a constrained lowpass filter, the backpanel and the frontpanel are bounded, from below and from above respectively, by the square root of the input image as shown in (6.6). These constraints can be relaxed in order to exploit some further degrees of freedom without compromising the feasibility of the solution.

Finally, by using a constrained lowpass filter to blur the backpanel, it is guaranteed that the frontpanel image will not exceed the white level. An examination of Figure 6.9, however, shows that the frontpanel computed in this way presents some undershoots near sharp edges; in some cases, these pixels may fall below the black level and dark details will be lost. In order to prevent also this artifact, some sort of double constraint must be used.

In order to derive the exact constraints, we shall suppose that each pixel of the two panels can have a normalized transmittance between $1/d$ and 1, where d is the dynamic range. The maximum transmittance of an LCD panel is actually much lower (0.3 for the IDTech grayscale panels used in the prototypes, 0.1 for color panels), but we will use normalized units for simplicity of notation. By definition, the backpanel does not saturate if the transmittance $T_{\mathrm{b}}(x, y)$ is within the limits:

$$\frac{1}{d} \leq T_{\mathrm{b}}(x, y) \leq 1 \qquad \forall (x, y) \tag{6.8}$$

If perfect reconstruction is requested, the frontpanel is computed by division. The constraints in this case are

$$\frac{1}{d} \leq \frac{L_{\text{in}}(x,y)}{B\,T_{\text{b}}(x,y)} \leq 1 \qquad \Rightarrow \qquad \frac{L_{\text{in}}(x,y)}{B} \leq T_{\text{b}}(x,y) \leq \frac{d\,L_{\text{in}}(x,y)}{B} \qquad \forall (x,y) \tag{6.9}$$

Therefore, merging the two sets of constraints (6.8) and (6.9), we obtain that both panels do not saturate if the backpanel satisfies

$$\max\left\{\frac{L_{\text{in}}(x,y)}{B}, \frac{1}{d}\right\} \leq T_{\text{b}}(x,y) \leq \min\left\{\frac{d\,L_{\text{in}}(x,y)}{B}, 1\right\} \qquad \forall (x,y) \tag{6.10}$$

If the dynamic range of the input image is greater than d^2, the constraints (6.10) become incompatible, therefore the input image should be pre-processed in order to remove any exceedingly dark pixels. This is unnecessary if the digital data of the source image are decoded following a DICOM grayscale display function which matches the dynamic range of the display device.

Based on the assumption that the parallax error is small if the backpanel image is smooth, we use an optimization procedure to generate the smoothest possible backpanel image subject to the constraints (6.10). Smoothness can be measured in several ways, and the measure we propose is suggested by the expression of the relative error (6.3). Since the backpanel we are seeking is a smooth function and the displacements Δx and Δy are small, we approximate (6.3) with a first-order Taylor expansion:

$$\begin{aligned} E(x,y) &\triangleq \frac{T_{\text{b}}(x+\Delta x, y+\Delta y) - T_{\text{b}}(x,y)}{T_{\text{b}}(x,y)} \approx \\ &\approx \frac{1}{T_{\text{b}}(x,y)}\left[\frac{\partial T_{\text{b}}(x,y)}{\partial x}\,\Delta x + \frac{\partial T_{\text{b}}(x,y)}{\partial y}\,\Delta y\right] = \frac{\nabla T_{\text{b}}(x,y)}{T_{\text{b}}(x,y)} \cdot \begin{bmatrix}\Delta x \\ \Delta y\end{bmatrix}, \end{aligned} \tag{6.11}$$

where $\cdot$ indicates a scalar product. By the Schwarz inequality, two arbitrary vectors $\mathbf{a}$ and $\mathbf{b}$ satisfy $|\mathbf{a}\cdot\mathbf{b}| \leq \|\mathbf{a}\|\,\|\mathbf{b}\|$, with the equality holding when the two vectors are parallel. Following this property, we choose to compute the mean square norm of the gradient term in the right hand side of Equation (6.11), in order to obtain a single scalar which can be minimized by an optimization procedure. The expression we obtain is the following one, in which the identity comes from simple calculus:

$$E_{\text{mean}} \triangleq \iint \left\| \frac{\nabla T_{\text{b}}(x,y)}{T_{\text{b}}(x,y)} \right\|^2 dx\,dy = \iint \left\| \nabla \log T_{\text{b}}(x,y) \right\|^2 dx\,dy. \tag{6.12}$$

This approach is different from the one described in Section 6.5: the backpanel is not computed by means of some sort of lowpass filter, applied to an appropriate signal such as the square root of the input image; it is instead generated by an optimization procedure and its shape is only determined by the constraints. The proposed method theoretically produces an optimal result, since the objective function (6.12) measures the parallax error and the constraints (6.10) are strictly derived from the device specifications. The method is also easy to tune because it does not contain any free parameter which requires a manual adjustment. An example of its behavior on a 1D test signal is shown in Figure 6.12. The algorithm produces a constant backpanel wherever possible; if the dynamic range of the input signal is less than d, the signal is displayed on the frontpanel only and the parallax error is completely avoided. In the bright areas, the

lower bound is active and the backpanel takes a value $T_{\mathrm{b}}(x, y) \geq L_{\mathrm{in}}(x, y)/B$; in the dark areas, instead, the upper bound is active and the backpanel takes a value $T_{\mathrm{b}}(x, y) \leq d\, L_{\mathrm{in}}(x, y)/B$. The step in the backpanel is therefore approximately d times smaller than the step in the input signal; it can also be noticed that the lower bound in (6.10) is lower than the square root of the input image, so there is a better chance that the algorithm produces a smooth backpanel. If no constraint is active, the backpanel forms a linear slope (in logarithmic scale) in order to minimize the functional (6.12).

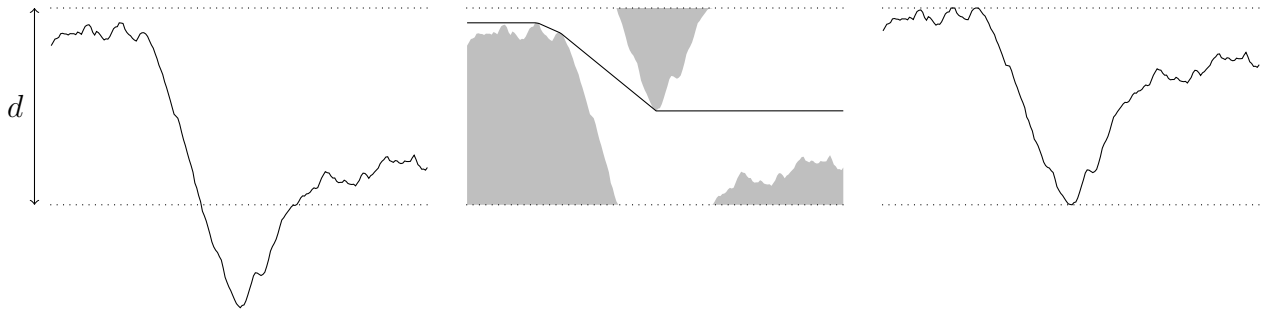

Figure 6.12: Second proposed method applied to a 1D test signal. Left to right: input signal, backpanel and constraints (shaded), frontpanel. The vertical axis is in logarithmic scale for a better legibility

On 2D images, the algorithm performs in a similar way. The backpanel computed with the proposed method is similar to a membrane-based envelope, with the difference that the membrane is not attracted to an input signal (mathematically, this means setting $\lambda = 0$ in (6.7)) and its shape is only determined by the constraints. In the areas where no constraint is active, the backpanel satisfies the Laplace equation. The use of higher-order smoothness functionals, such as a thin-plate-spline bending energy, was not considered both for computational efficiency and for a theoretical motivation. The membrane smoothness functional minimizes the average slope of the backpanel, and therefore produces a constant output when possible; the spline functional instead minimizes the bending energy, and does not penalize a steep gradient.

Performing a constrained optimization with a large number of unknowns might seem a computationally very complex task. On the contrary, the proposed method can be implemented very efficiently by means of multigrid techniques [12, 32]. An unoptimized C implementation runs in approximately 0.4 seconds on a 1.5 GHz Pentium 4 processor for a 1 megapixel image, and the running time is approximately proportional to the number of pixels in the image. It can be noticed that, if the dynamic range of the input image is lower that the dynamic range d of the panels, the optimization problem admits multiple solutions, with the backpanel equal to a constant value and the frontpanel proportional to the input image. This is not an issue in practice, because the different solutions produce the same visual output and can therefore be chosen arbitrarily. Moreover, in a practical implementation, the optimization problem is solved numerically by means of an iterative method, which converges to just one solution depending on the initialization. A modification in the objective function (6.12), in order to guarantee a unique solution, would actually increase the algorithm complexity without introducing a practical benefit. Moreover, this issue can not occur if the input image is encoded following the DICOM grayscale display function [52]. In this case, the lowest and highest digital levels of the encoded image are mapped to the black and white levels of the display respectively, and the dynamic range of the input image has therefore a fixed value.

6.7 Method 3: Parallax error reduction with loss of perfect reconstruction

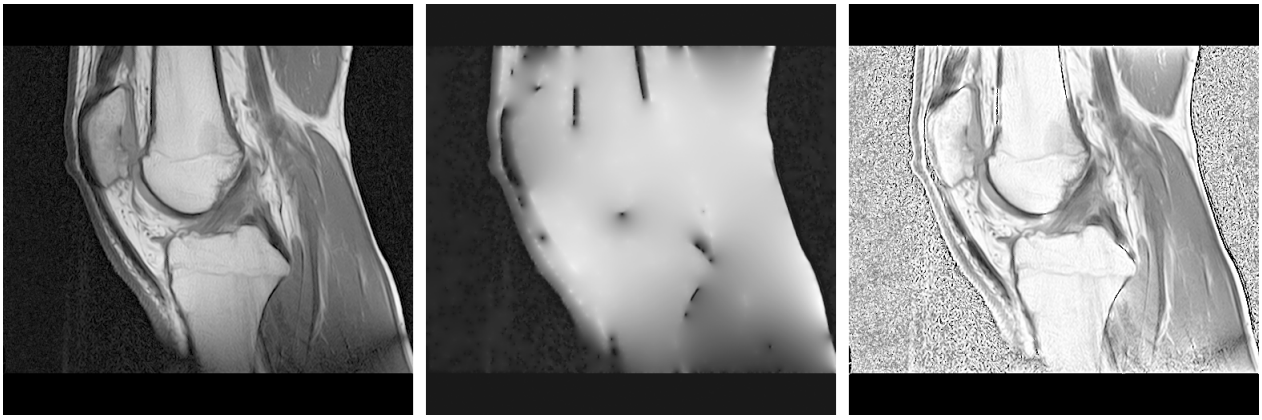

Figure 6.13: Example of splitting. Left to right: original image, backpanel, frontpanel.

The algorithm described in Section 6.6 can not eliminate completely the parallax effects if the input image contains sharp edges which have a greater magnitude than the dynamic range d of the panels. In this case, due to the presence of the upper bound in (6.10), the frontpanel alone is not able to completely reproduce the edge, and a fraction of its magnitude must be transferred onto the backpanel. An example is visible in Figure 6.13: the source image contains sharp edges and isolated black pixels; if perfect reconstruction is requested, a very dark pixel can only be reproduced when both panels are dark and the degrees of freedom allowed by the constraints are strongly reduced. Similar issues are also present in other HDR display technologies [67] and are currently unavoidable.

Since the upper bound was introduced to guarantee perfect reconstruction and detail preservation, we must deduce that these requirements are incompatible with that of parallax error reduction. A possible solution to this problem is to relax the perfect reconstruction constraint (6.4) and allow some distortion in the visualized image. By observing Figure 6.13, it can be noticed that the dark pixels typically carry little information content; if this were true in general, one could filter the input image before the splitting in order to remove isolated black pixels and limit sharp edges. It is also known [51, 77] that the high dynamic range of the human visual system is mainly due to local adaptation, and if a scene contains high contrast boundaries the details near the edge appear blurred and indistinct. Because of this limitation, in some cases a reconstruction error near a sharp edge might not be visible.

6.7.1 Variational formulation

In order to relax the prefect reconstruction constraint, we propose [25] to formulate the splitting task as an optimization problem. This technique may appear computationally expensive, since it involves the minimization of a functional with several million unknowns. For each input image, two output images must be computed for the backpanel and the frontpanel respectively, and since a constrained optimization is being performed in order to prevent any clipping, extra variables may be added in the form of Lagrange multipliers. However, we were able to derive

an objective function that can be minimized very efficiently by means of multigrid methods [32], as shown in the following.

The image splitting for a dual layer display without the perfect reconstruction constraint is a multi-objective optimization problem, in which both the parallax error and the reconstruction error must be minimized. The problem can be formally stated by defining two functionals E_{par} and E_{rec} that measure these two errors respectively. The unknown functions are the backpanel transmittance $T_{\text{b}}(x, y)$ and frontpanel transmittance $T_{\text{f}}(x, y)$, where (x, y) are the pixel coordinates. Due to the loss of perfect reconstruction, the frontpanel image can not be computed by division and must be treated as a second unknown. In order to prevent clipping, $T_{\text{b}}(x, y)$ and $T_{\text{f}}(x, y)$ are subject to simple bound constraints, which are constant over the entire image:

$$\begin{array}{l} T_{\min} \leq T_{\text{b}}(x, y) \leq T_{\max} \\ T_{\min} \leq T_{\text{f}}(x, y) \leq T_{\max} \end{array} \qquad \forall (x, y), \tag{6.13}$$

where $T_{\min}$ and $T_{\max}$ are the black and white levels of the panels.

A separate minimization of the two functionals does not provide any useful solution. The parallax error is clearly minimized when the backpanel takes an arbitrary constant value, but a constant backpanel of course nullifies the benefit of the dual layer technology. The reconstruction error is minimized when perfect reconstruction is achieved, but the perfect reconstruction constraint inevitably produces an irregular backpanel when the input image contains sharp edges. The two objectives are mutually incompatible and some sort of tradeoff must be found. We minimize a weighted sum of the two functionals

$$E_{\text{total}} \triangleq k_{\text{p}}\, E_{\text{par}} + E_{\text{rec}} = \text{Min}, \tag{6.14}$$

where k_{p} is a user-adjustable scalar parameter. In the rest of this section we will assume that $T_{\text{b}}(x, y)$ and $T_{\text{f}}(x, y)$ are functions of continuous variables. Besides simplifying the notation, this allows us to use the calculus of variations [33] in order to easily formulate the problem and derive the corresponding equations. The practical implementation of the proposed method involves an appropriate discretization of the differential equations and the use of a numerical method.

6.7.2 Parallax error metric

Following the approach used in Section 6.6, we propose to measure the parallax error with the functional

$$E_{\text{par}} \triangleq \iint \frac{1}{2} \left\| \frac{\nabla T_{\text{b}}(x, y)}{T_{\text{b}}(x, y)} \right\|^2 dx\, dy. \tag{6.15}$$

The use of a ratio instead of a gradient alone accounts for Weber's law and provides a more perceptually significant smoothness measure. The integral is computed over the whole image area, and the factor $1/2$ is introduced in order to simplify future calculations. From simple calculus, it follows that (6.15) is equivalent to

$$E_{\text{par}} \triangleq \iint \frac{1}{2} \left\| \nabla \log(T_{\text{b}}) \right\|^2 dx\, dy, \tag{6.16}$$

which allows to avoid the division and to obtain a simple quadratic functional by means of a logarithmic change of variables.

Since Weber's law was shown to hold only for high luminance levels, a more accurate definition should ideally take into account the loss of sensitivity in dark areas. However, the backpanel transmittance is not simply related to the output luminance due to the presence of the frontpanel. In order to build a more accurate model, the functional E_{par} should be made dependent on the frontpanel image too, but the added complexity is unlikely to produce a substantial quality improvement. Indeed, Weber's law overestimates the eye sensitivity in dark areas, therefore the minimization of the functional (6.16) is likely to produce an even smoother backpanel.

6.7.3 Reconstruction error metric

The definition of the reconstruction error metric E_{rec} poses the most difficult issues, because it must satisfy two strongly conflicting requirements. Its expression should be kept as simple as possible, because it is being used inside an optimization algorithm; at the same time, it should be accurate enough to allow a substantial relaxation of the perfect reconstruction constraints, which is necessary to produce a smoother backpanel also in correspondence of sharp edges in the input image. Accurate visible difference predictors such as [15] or [47] have a very high computational complexity, and can only be used in practice to evaluate *a posteriori* the fidelity of two given images or to perform the unsupervised tuning of a limited set of parameters. Attempting to generate an entire image by means of an optimization procedure would involve extremely long computational times, thus making the system useless in practice. Moreover, due to the high complexity and nonlinear nature of the models, it is difficult to verify whether the solution found by the optimization algorithm is indeed a global minimum. In order to make the algorithm feasible in practice, we propose to use a simplified visible difference metric that considers only a reduced set of phenomena. Besides simplifying the computation, this approach might actually produce a more accurate reconstruction. Indeed, many stages of the human visual system have the effect of reducing the visibility of a certain class of details; if these parts are neglected, the model is going to overestimate the visibility of artifacts, and as a consequence the minimization of E_{par} is going to produce a smaller reconstruction error. The tradeoff, as previously noticed, is that a more faithful reconstruction reduces the degrees of freedom available for the parallax error reduction and therefore produces a more irregular backpanel.

In the definition of the reconstruction error metric E_{rec} we pay attention primarily to the computational aspects, and verify experimentally if the results are satisfactory. Simple distance metrics that only consider the pointwise error, such as for instance the familiar Mean Squared Error (MSE) or Peak Signal to Noise Ratio (PSNR), often provide a poor estimate because the visibility of an error strongly depends on its spatial distribution and on the background. The frequency selectivity of the first stages of the HVS (optical blurring, photoreceptor sampling and retinal cell interactions) is described by the contrast sensitivity function (CSF) [7]. However, distance metrics that take into account the CSF involve a spatial filtering, and are already very difficult to incorporate into an optimization algorithm, since their minimization would involve some sort of deconvolution process. Methods based on the Discrete Fourier Transform (DFT) can not be used, because a *constrained* optimization problem is being solved, and constraints expressed in the space domain can not be easily mapped to the frequency domain. Moreover, the human contrast sensitivity depends on the background luminance; as a consequence, a rigorous model can not be based on a straightforward linear convolution and methods based on linear

transforms lose their efficacy [47]. This is especially critical when HDR images are considered such as in the present application. Finally, the computation of the CSF requires the accurate knowledge of the observer's distance and position. Daly [15], for instance, assumes that the viewing distance falls within a determined range, and considers an envelope of the corresponding CSFs. Once the CSF is neglected, we are forced to neglect also the phenomena, such as visual masking, which take place at a later stage inside the visual cortex. Such phenomena are not yet fully understood, and the models that have been proposed so far have a high complexity.

The simplified metric we propose only takes into account the phenomenon of *local adaptation*. The human eye is able to perceive brightness values in a range of over 10 orders of magnitude from starlight to sunlight, and at the same time to detect fine local variations. The photoreceptors in the retina achieve this ability by modulating their response as a function of the local mean luminance, or *adaptation level*. The eye sensitivity can be quantified using the concept of *just noticeable difference* (JND), which represents the minimum luminance variation ΔL an average observer can detect at a given adaptation level L_{ad}. The JND can be measured psychophysically in *detection threshold* experiments [20], and was shown to increase with the adaptation level, thus implying that the same luminance difference can have a radically different visibility depending on its surroundings. For a wide range of luminances the relation is well approximated by a direct proportionality; this rule is known as the Weber Law or as the Weber-Fechner law, and holds for other kinds of sensorial stimuli as well. For very dark or very bright background levels, saturation mechanisms occur and this simplified law does not hold any more; in general, the relationship between the JND and the background luminance follows a curve known as the *threshold versus intensity* (TVI) function [1]: $\Delta L = \mathrm{TVI}(L_{\mathrm{ad}})$. We choose to estimate the visibility of the error by comparing the pointwise difference between the input image $L_{\mathrm{in}}(x, y)$ and the output image $L_{\mathrm{out}}(x, y) = B\, T_{\mathrm{b}}(x, y)\, T_{\mathrm{f}}(x, y)$ with the JND computed in the corresponding point. If we assume that the image falls undistorted onto the retina, that the cone cells are able to adapt to the single pixel, and that the JND can be computed using Weber's law, we obtain that the pointwise perceptual distance between the original image $L_{\mathrm{in}}(x, y)$ and the reconstructed image $L_{\mathrm{out}}(x, y)$ can be estimated using the relative error:

$$E(x, y) \triangleq \frac{L_{\mathrm{out}}(x, y) - L_{\mathrm{in}}(x, y)}{L_{\mathrm{in}}(x, y)} = \frac{B\, T_{\mathrm{b}}(x, y)\, T_{\mathrm{f}}(x, y) - L_{\mathrm{in}}(x, y)}{L_{\mathrm{in}}(x, y)}. \tag{6.17}$$

In reality the eye can not adapt to the single pixel, and the local adaptation can be easily taken into account by replacing the relative error in Eq. (6.17) with the modified expression

$$E_{\mathrm{ad}}(x, y) \triangleq \frac{B\, T_{\mathrm{b}}(x, y)\, T_{\mathrm{f}}(x, y) - L_{\mathrm{in}}(x, y)}{L_{\mathrm{ad}}(x, y)}, \tag{6.18}$$

where L_{ad} is the eye adaptation level. Since L_{ad} corresponds to some sort of local mean luminance, in dark pixels close to a bright area we have $L_{\mathrm{ad}}(x, y) > L_{\mathrm{in}}(x, y)$, therefore the visual threshold increases and our model predicts the loss of sensitivity that is commonly experienced in these cases. An accurate computation of L_{ad} is still an open research subject; in the following we will only focus on computing T_{b} and T_{f} assuming L_{ad} is given. The Weber law failure for low luminance levels can be easily taken into account by replacing $L_{\mathrm{ad}}(x, y)$ in the denominator of (6.18) with $\mathrm{TVI}\big(L_{\mathrm{ad}}(x, y)\big)$.

Equation (6.18) gives an estimate of the pointwise perceptual distance between the input and the reconstructed image. The functional E_{rec} is obtained by computing the quadratic sum of

(6.18)

$$E_{\text{rec}} \triangleq \iint \frac{1}{2}\left[\frac{B\,T_{\text{b}}(x,y)\,T_{\text{f}}(x,y) - L_{\text{in}}(x,y)}{\text{TVI}\big(L_{\text{ad}}(x,y)\big)}\right]^2 dx\,dy. \tag{6.19}$$

As in (6.16), the integral is computed over the entire image, and the factor $1/2$ is introduced in order to simplify future calculations. The use of a quadratic sum is motivated by computational considerations, since it will give rise to a quadratic functional after an appropriate change of variables described in the following.

From this point on, we will drop the indices (x, y) for simplicity of notation. We have shown in Section 6.7.2 that E_{par} becomes a quadratic functional if a simple logarithmic transformation is performed. A similar simplification is possible also for E_{rec}, and uses the following approximation, based on a first-order Taylor expansion:

$$\begin{aligned} &\log(L_{\text{out}}) - \log(L_{\text{in}}) = \log(L_{\text{in}} + L_{\text{out}} - L_{\text{in}}) - \log(L_{\text{in}}) = \\ &= \log(L_{\text{in}}) + \frac{L_{\text{out}} - L_{\text{in}}}{L_{\text{in}}} + o(L_{\text{out}} - L_{\text{in}}) - \log(L_{\text{in}}) \approx \frac{L_{\text{out}} - L_{\text{in}}}{L_{\text{in}}}, \end{aligned} \tag{6.20}$$

where $o(L_{\text{out}} - L_{\text{in}})$ represents the higher-order terms. The presence of dimensional variables such as luminances inside a logarithm does not introduce an inconsistency in this case, because the expression (6.20) contains logarithmic differences which cancel out any scale factors; the expression above is therefore adimensional. By using the derived approximation, the expression (6.18) can be re-written as:

$$\begin{aligned} &\frac{B\,T_{\text{b}}\,T_{\text{f}} - L_{\text{in}}}{\text{TVI}(L_{\text{ad}})} = \frac{L_{\text{out}} - L_{\text{in}}}{\text{TVI}(L_{\text{ad}})} = \frac{L_{\text{in}}}{\text{TVI}(L_{\text{ad}})}\frac{L_{\text{out}} - L_{\text{in}}}{L_{\text{in}}} \approx \\ &\approx \frac{L_{\text{in}}}{\text{TVI}(L_{\text{ad}})}\big[\log(L_{\text{out}}) - \log(L_{\text{in}})\big] = \frac{L_{\text{in}}}{\text{TVI}(L_{\text{ad}})}\big[\log(T_{\text{b}}) + \log(T_{\text{f}}) - \log(L_{\text{in}}/B)\big]. \end{aligned} \tag{6.21}$$

It can be noticed that the use of an adaptive nonlinearity in (6.18) compensates the absence of the CSF in the proposed model. Indeed, the CSF has a bandpass response, in which the low-frequency drop-off is due to the local adaptation. Models such as [15, 47] contain the CSF, but only use a simple point nonlinearity as a front-end. Our formulation brings significant advantages in terms of computational efficiency, because L_{ad} is computed as a function of the input luminance, and is a known term in the optimization problem. Improved models for the computation of L_{ad} can be readily incorporated into the image splitting algorithm without altering its structure and complexity.

6.7.4 Global functional

By combining the two functionals E_{par} and E_{rec} as in (6.14), the functional to be minimized takes the following form:

$$E_{\text{total}} = \iint \left\{ \frac{k_{\text{p}}}{2}\big\|\nabla \log(T_{\text{b}})\big\|^2 + \frac{1}{2}\left(\frac{L_{\text{in}}}{\text{TVI}(L_{\text{ad}})}\right)^2 \big[\log(T_{\text{b}}) + \log(T_{\text{f}}) - \log(L_{\text{in}}/B)\big]^2 \right\} dx\,dy. \tag{6.22}$$

For simplicity of notation, we define

$$w_{\text{ad}} \triangleq \left(\frac{L_{\text{in}}}{\text{TVI}(L_{\text{ad}})}\right)^2.$$

As previously noticed, in dark areas close to bright areas we have $L_{\text{ad}} > L_{\text{in}}$, therefore w_{ad} takes a small value and in those points the optimization procedure privileges a smooth backpanel at the cost of some error in the reconstructed image. By making the change of variables $\mathcal{T}_{\text{b}} \triangleq \log(T_{\text{b}})$ and $\mathcal{T}_{\text{f}} \triangleq \log(T_{\text{f}})$, (6.22) becomes a quadratic functional in the new unknowns $\mathcal{T}_{\text{b}}$ and $\mathcal{T}_{\text{f}}$:

$$\iint \left\{ \frac{k_{\text{p}}}{2} \left\| \nabla \mathcal{T}_{\text{b}} \right\|^2 + \frac{w_{\text{ad}}}{2} \left[\mathcal{T}_{\text{b}} + \mathcal{T}_{\text{f}} - \log(L_{\text{in}}/B) \right]^2 \right\} dx\, dy, \tag{6.23}$$

subject to the constraints

$$\begin{array}{l} \log(T_{\min}) \leq \mathcal{T}_{\text{b}}(x,y) \leq \log(T_{\max}) \\ \log(T_{\min}) \leq \mathcal{T}_{\text{f}}(x,y) \leq \log(T_{\max}) \end{array} \qquad \forall (x,y). \tag{6.24}$$

The constraints are incorporated by adding two Lagrange multipliers λ and μ [9]:

$$\iint \left\{ \frac{k_{\text{p}}}{2} \left\| \nabla \mathcal{T}_{\text{b}} \right\|^2 + \frac{w_{\text{ad}}}{2} \left[\mathcal{T}_{\text{b}} + \mathcal{T}_{\text{f}} - \log(L_{\text{in}}/B) \right]^2 + \lambda \mathcal{T}_{\text{b}} + \mu \mathcal{T}_{\text{f}} \right\} dx\, dy. \tag{6.25}$$

From the Karush-Kuhn-Tucker theorem of nonlinear programming, it follows that the Lagrange multipliers meet the conditions listed in Table 6.1. In particular, the Lagrange multipliers — which are functions of (x, y) — are nonzero only in the points where the corresponding constraint is active.

Value of $\mathcal{T}_{\text{b}}$	Sign of λ
$\mathcal{T}_{\text{b}} = \log(T_{\min})$	$\lambda < 0$
$\log(T_{\min}) < \mathcal{T}_{\text{b}} < \log(T_{\max})$	$\lambda = 0$
$\mathcal{T}_{\text{b}} = \log(T_{\max})$	$\lambda > 0$

Value of $\mathcal{T}_{\text{f}}$	Sign of μ
$\mathcal{T}_{\text{f}} = \log(T_{\min})$	$\mu < 0$
$\log(T_{\min}) < \mathcal{T}_{\text{f}} < \log(T_{\max})$	$\mu = 0$
$\mathcal{T}_{\text{f}} = \log(T_{\max})$	$\mu > 0$

Table 6.1: Sign of the Lagrange multipliers in Equation (6.25)

The Euler-Lagrange differential equation associated to the functional (6.25) is

$$\left\{ \begin{array}{l} w_{\text{ad}} \left[\mathcal{T}_{\text{b}} + \mathcal{T}_{\text{f}} - \log(L_{\text{in}}/B) \right] + \lambda = k_{\text{p}} \nabla^2 \mathcal{T}_{\text{b}} \\ w_{\text{ad}} \left[\mathcal{T}_{\text{b}} + \mathcal{T}_{\text{f}} - \log(L_{\text{in}}/B) \right] + \mu = 0 \end{array} \right. \tag{6.26}$$

The factors 1/2 which were introduced in (6.15) and (6.19) remove a factor 2 that would otherwise be present in (6.26).

In the points where no constraint is active, $\mathcal{T}_{\text{b}}$ satisfies a Laplace equation and perfect reconstruction is achieved. In fact, setting both λ and μ to zero and simplifying, the equations (6.26) reduce to

$$\left\{ \begin{array}{l} \nabla^2 \mathcal{T}_{\text{b}} = 0 \\ \mathcal{T}_{\text{b}} + \mathcal{T}_{\text{f}} = \log(L_{\text{in}}/B) \end{array} \right. \tag{6.27}$$

In the points where the backpanel is constrained and the frontpanel is free, we have

$$\left\{ \begin{array}{l} \mathcal{T}_{\text{b}} = \log(T_{\min}) \quad \text{or} \quad \mathcal{T}_{\text{b}} = \log(T_{\max}) \\ \mathcal{T}_{\text{b}} + \mathcal{T}_{\text{f}} = \log(L_{\text{in}}/B) \end{array} \right. \tag{6.28}$$

Again, perfect reconstruction is achieved. However, we shall prove in the following Section 6.8 that this case does not occur in practice, and the knowledge of this property can be used to

improve the performance of the computational algorithm by avoiding unnecessary tests. Perfect reconstruction is lost only in the points where the frontpanel hits a constraint; in this case we have

$$\begin{cases} w_{\text{ad}} \left[\mathcal{T}_{\text{b}} + \mathcal{T}_{\text{f}} - \log(L_{\text{in}}/B)\right] = k_{\text{p}} \nabla^2 \mathcal{T}_{\text{b}} \\ \mathcal{T}_{\text{f}} = \log(T_{\min}) \quad \text{or} \quad \mathcal{T}_{\text{f}} = \log(T_{\max}) \end{cases} \tag{6.29}$$

6.8 Constraint simplification

If the method proposed in Section 6.7 is used, one computational simplification is obtained by noticing that the constraints on the backpanel are never active and can therefore be dropped [26]. Let us suppose for instance that the backpanel is constrained to its upper boundary, i.e. $\mathcal{T}_{\text{b}}(x, y) = \log(T_{\max})$ in some point (x, y), while the frontpanel is free in the same point. As previously shown, this implies that the Lagrange multipliers satisfy $\lambda(x, y) > 0$ and $\mu(x, y) = 0$. As a consequence, the second equation in (6.26) shows that the term in square brackets is zero, and the first equation becomes $k_{\text{p}} \nabla^2 \mathcal{T}_{\text{b}}(x, y) = \lambda(x, y) > 0$. However, a positive Laplacian indicates that the function $\mathcal{T}_{\text{b}}$ has a "$\cup$" shape in the neighborhood of (x, y) and would therefore exceed the upper bound in some adjacent point. Following a similar reasoning, if the backpanel were constrained to its lower bound $\mathcal{T}_{\text{b}} = \log(T_{\min})$, it would necessarily have a "$\cap$" shape and violate the bound.

If we suppose that both panels are constrained to their upper bound, then under the reasonable assumption that the input image falls within the dynamic range of the dual layer display (i.e. $T_{\min}^2 \leq L_{\text{in}}(x, y)/B \leq T_{\max}^2 \ \forall (x, y)$) we obtain that the term in square brackets in Equation (6.26) is always ≥ 0. Since $\mu(x, y) > 0$, the second equation can not be satisfied; moreover, as in the previous case we would obtain $\nabla^2 \mathcal{T}_{\text{b}}(x, y) > 0$, causing the backpanel to violate the constraints. The symmetric case of both panels constrained to the lower bound is excluded following a similar reasoning.

Finally, if $\mathcal{T}_{\text{b}}(x, y) = \log(T_{\max})$ and $\mathcal{T}_{\text{f}}(x, y) = \log(T_{\min})$, the term in square brackets must be strictly positive in order to satisfy the second equation in (6.26), since $\mu(x, y) < 0$. This again would imply a positive Laplacian, because $\lambda(x, y) > 0$. In a similar way, the symmetric case $\mathcal{T}_{\text{b}}(x, y) = \log(T_{\min})$ and $\mathcal{T}_{\text{f}}(x, y) = \log(T_{\max})$ would also imply a violation of the constraints. We conclude therefore that the constraints on the backpanel are never active, provided that the input image falls within the display dynamic range as assumed.

6.9 Numerical method

Equation (6.25) was derived for functions of continuous variables for simplicity of notation. In a practical implementation the equations must be discretized, both because the images are sampled and because the equation can not be solved in closed form except in some very simple cases.

6.9.1 Solution of a simplified problem

In order to describe more clearly the proposed numerical method [26], we shall first describe its principles on a simplified version of the functional, with one unknown and no constraints, and then show how to extend the method to the present case, with two unknowns and constraints. The simplified functional we shall consider is:

$$E_{\text{cont}} \triangleq \iint \left\{ \frac{1}{2} \big\| \nabla f(x,y) \big\|^2 + \frac{w(x,y)}{2} \big[f(x,y) - u(x,y) \big]^2 \right\} dx\, dy, \tag{6.30}$$

where $f(x,y)$ is the unknown function, $u(x,y)$ is a known input term and $w(x,y)$ is a weight factor which can be space-dependent in general. The Euler-Lagrange differential equation associated to the functional (6.30) is a Helmholtz equation:

$$w(x,y) \big[f(x,y) - u(x,y) \big] = \nabla^2 f(x,y), \tag{6.31}$$

with the natural boundary condition $\nabla f(x,y) \cdot \hat{n} = 0$.

The choice of the discretization scheme follows naturally from the fact that the functions $f(x,y)$ and $u(x,y)$ are digital images, sampled on a uniform Cartesian grid, and are therefore represented by the two-dimensional arrays $f[i,j]$ and $u[i,j]$ of size $M \times N$. We shall assume for simplicity that the sampling has the same step size h in both horizontal and vertical directions. Rather than discretizing the differential equation (6.31), it is convenient to discretize the functional (6.30) by replacing the gradient with finite differences and the integral with a summation. In this way an algebraic function – more precisely, a quadratic form – is obtained. A possible discretization scheme is the following one, in which the limits of the summation are chosen appropriately in order to take into account the image boundary:

$$\begin{aligned} E_{\text{dis}} = {} & \frac{h^2}{2} \sum_{i=1}^{M} \sum_{j=1}^{N-1} \left(\frac{f[i,j+1] - f[i,j]}{h} \right)^2 + \\ & + \frac{h^2}{2} \sum_{i=1}^{M-1} \sum_{j=1}^{N} \left(\frac{f[i+1,j] - f[i,j]}{h} \right)^2 + \\ & + \frac{h^2}{2} \sum_{i=1}^{M} \sum_{j=1}^{N} w[i,j] \big(f[i,j] - u[i,j] \big)^2. \end{aligned} \tag{6.32}$$

By taking the partial derivatives of the function (6.32) with respect to each variable and setting them equal to zero, a linear system of size $MN \times MN$ is obtained. If we define, for simplicity of notation, $S[i,j] \triangleq f[i+1,j] + f[i-1,j] + f[i,j+1] + f[i,j-1]$, the equations we obtain have the form:

$$\frac{\partial E_{\text{dis}}}{\partial f[i,j]} = 4f[i,j] - S[i,j] + h^2\, w[i,j] \big(f[i,j] - u[i,j] \big) = 0 \tag{6.33}$$

The equations (6.33) hold also on the image boundary, provided that in the definition of $S[i,j]$ any pixel which falls outside the boundary is replaced with $f[i,j]$. Besides handling the boundary conditions in a natural way, another advantage obtained by discretizing the functional (6.30) rather than the differential equation (6.31) comes from the fact that it is generally easier to minimize a function than to solve an arbitrary linear system. For instance,

it is immediately clear that the linear system (6.33) is positive definite, since it is made of the partial derivatives of a quadratic form, which is itself positive definite being a sum of squares.

The numerical solution of (6.30) or (6.31) is obtained by solving the linear system (6.33). Since each unknown $f[i,j]$ represents one pixel of the image, for high resolution images the system may contain millions of unknowns, and direct methods such as the familiar Gaussian elimination, LU decomposition or Cholesky decomposition would require very high processing times. The system is however highly sparse and has a particular structure; therefore, iterative methods can be used profitably. One simple iterative method is *point-by-point minimization*, in which the quadratic form (6.32) is minimized by changing only one variable at a time. This method is equivalent to solving each equation of the linear system (6.33) in turn for the variable $f[i,j]$:

$$f[i,j] \leftarrow \frac{S[i,j] + h^2\, w[i,j]\, u[i,j]}{4 + h^2\, w[i,j]}. \tag{6.34}$$

The point-by-point minimization is therefore equivalent to the Gauss-Seidel iterative method [64]. In this way, at each iteration the value of the objective function decreases (since a minimization is being performed) and the solution is approached. It can be noticed that, for boundary pixels, the $S[i,j]$ term also depends on $f[i,j]$; this however does not hinder the convergence and simplifies the algorithm by allowing the same expression (6.34) for all the pixels.

A Gauss-Seidel iteration involves very simple operations and uses a limited amount of memory, because the variables are updated "in place", but can not be parallelized in general because, due to the presence of the $S[i,j]$ term in (6.34), the updated value of each pixel must be available before the adjacent ones can be processed. One alternative method is the Jacobi method [64], in which each variable $f[i,j]$ is updated using the values from the previous iteration. In this way parallel processing is possible, because the dependencies between adjacent pixels are removed, but the memory requirements are increased because a temporary copy of $f[i,j]$ must be stored. If the Gauss-Seidel method is used, the dependencies can be avoided by simply processing the pixels in a different order. A popular scheme is the *red-black* one, in which the pixels are divided into two groups resembling the two colors of a chessboard. It is easy to see that, when for instance a "red" pixel is processed, the right hand side of expression (6.34) only contains "black" pixels; it is therefore possible to process all "red" pixels in parallel and the point-by-point minimization is actually equivalent to minimizing the function (6.33) for half the variables simultaneously. In a similar way, the "black" pixels can also be processed simultaneously. In modern processors, removing the dependencies between instructions can remove performance bottlenecks thanks to the presence of the pipeline, without the need of modifying the code. The performance can be further enhanced by using SIMD (Single Instruction, Multiple Data) instructions which explicitly perform two or four operations simultaneously, provided that the software is appropriately coded. In this case, it may be convenient to shuffle the pixels in the memory in such a way that the pixels to be processed in parallel are stored in contiguous memory locations.

6.9.2 Overview of multigrid methods

Simple iterative methods such as the cited Gauss-Seidel or Jacobi have a poor performance when used alone. It is easy to see from (6.34) that, at each iteration, the information is propagated by one pixel, and this causes the low frequency components of the solution to

converge very slowly. Low frequencies, however, can be sampled on a coarse grid according to the Nyquist theorem. The basic idea of multigrid methods [32] consists in generating an approximate solution by downsampling the input image (typically by a factor 2), processing it and interpolating the result back to the original resolution. A few Gauss-Seidel iterations are then enough to reconstruct also the high frequencies. Low frequencies in the original image correspond to higher frequencies in the downsampled image, and the convergence is therefore accelerated. The low resolution image can be processed recursively in the same way until the number of unknowns becomes small enough to permit a direct solution of the equation.

The practical implementation of multigrid methods can pose some difficulties. For instance, the low-resolution problem is described by a different linear system and the transformation is not straightforward, in particular when the equation contains spatially-varying coefficients such as in the present case. One common approach was introduced by Galerkin, and is well known in engineering because it is used in the finite element method to approximate an infinite-dimensional problem (a partial differential equation) with a finite-dimensional one (an algebraic system). It can be described in several different ways; we chose the following one, which gives an insight of the theoretical motivation underlying the method.

We wish to minimize a quadratic form, such as (6.32), and the exact solution is obtained by solving the linear system (6.33). In order to reduce the number of unknowns, the solution $f[i,j]$ is approximated by a linear combination of basis functions $\phi_k[i,j]$:

$$\tilde{f}[i,j] \triangleq \sum_{k=1}^{K} c_k \phi_k[i,j], \tag{6.35}$$

with $K < MN$. The coefficients c_k which provide the "best" approximation are computed by replacing the unknown array $f[i,j]$ in (6.32) with the approximate expression $\tilde{f}[i,j]$ and minimizing the resulting quadratic form with respect to the new unknowns c_k. This transformation, known as *Galerkin projection*, produces a linear system in K unknowns; once this reduced system is solved, the approximate solution $\tilde{f}[i,j]$ is computed using (6.35). In the present application, the coefficients c_k represent the pixel values of the downsampled image and the basis functions $\phi_k[i,j]$ perform an interpolation. In general, the reduced system generated by means of the Galerkin projection has a different form compared to the original one (6.33). It is possible to show (the computation is quite tedious and is omitted here) that, if the interpolation is performed by simply replicating each pixel of the image four times (or, in other words, if a nearest-neighbor interpolation by a factor 2 is used), the reduced system has the same form, up to a scale factor, as the original one. In this way, the equations at each resolution level are generated in a straightforward manner and the iterative refinement is always performed using the method described in the previous subsection. In our tests, we used this technique.

6.9.3 Solution of the actual optimization problem

The method described in the previous subsections – based on a point-by-point minimization scheme (or Gauss-Seidel iteration) combined with a multigrid method to improve the convergence speed – can be adapted with a few modifications to the present problem, in which two unknown functions $\mathcal{T}_\mathrm{b}$ and $\mathcal{T}_\mathrm{f}$ are involved and constraints are present on the frontpanel. It can be noticed that, when two variables are involved, the quadratic form obtained by discretizing (6.23) becomes positive semidefinite and in general multiple solutions are possible. Indeed, if

a constant value is added to $\mathcal{T}_\mathrm{b}$ and the same value is subtracted from $\mathcal{T}_\mathrm{f}$, the value of the objective function does not change. In this case, the constraints play a fundamental role in reducing the number of possible solutions. The problem can still admit multiple solutions in some cases, in particular when the input image has a low dynamic range and the constraints are not active, but this is not an issue in practice because the different solutions produce the same result and can therefore be chosen arbitrarily.

The variational problem (6.25) is discretized following the approach described in Subsection 6.9.1. The numerical method we propose is based on a modified point-by-point minimization scheme in which, for each pixel $[i, j]$ in turn, a constrained minimization of the two variables $\mathcal{T}_\mathrm{b}[i, j]$ and $\mathcal{T}_\mathrm{f}[i, j]$ is performed. We set, for simplicity of notation, $S[i, j] \triangleq \mathcal{T}_\mathrm{b}[i-1, j] + \mathcal{T}_\mathrm{b}[i+1, j] + \mathcal{T}_\mathrm{b}[i, j-1] + \mathcal{T}_\mathrm{b}[i, j+1]$, with the convention that the pixels which fall outside the image boundary are replaced with $\mathcal{T}_\mathrm{b}[i, j]$ as in Section 6.9.1. As a first step, a trial solution is computed by performing an unconstrained optimization of the discretized functional with respect to $\mathcal{T}_\mathrm{b}[i, j]$ and $\mathcal{T}_\mathrm{f}[i, j]$. This corresponds to solving the linear system

$$\begin{cases} h^2\, w_\mathrm{ad}[i, j]\big(\mathcal{T}_\mathrm{b}[i, j] + \mathcal{T}_\mathrm{f}[i, j] - \log(L_\mathrm{in}[i, j]/B)\big) = k_\mathrm{p}\big(S[i, j] - 4\mathcal{T}_\mathrm{b}[i, j]\big) \\ h^2\, w_\mathrm{ad}[i, j]\big(\mathcal{T}_\mathrm{b}[i, j] + \mathcal{T}_\mathrm{f}[i, j] - \log(L_\mathrm{in}[i, j]/B)\big) = 0 \end{cases} \tag{6.36}$$

The solution is

$$\mathcal{T}_\mathrm{b}[i, j] \leftarrow \frac{1}{4} S[i, j] \qquad \mathcal{T}_\mathrm{f}[i, j] \leftarrow \log(L_\mathrm{in}[i, j]/B) - \mathcal{T}_\mathrm{b}[i, j] \tag{6.37}$$

The effect of this iteration is to smooth the backpanel, by replacing the pixel $\mathcal{T}_\mathrm{b}[i, j]$ with an average of its four neighbors, and to update the frontpanel in order to guarantee perfect reconstruction, which as shown in Section 6.7 is still achieved in the points where no constraint is active. If the frontpanel $\mathcal{T}_\mathrm{f}[i, j]$ computed in this way satisfies the bounds, the trial solution (6.37) is accepted and the next pixel is processed; if instead the frontpanel exceeds one bound, a constrained minimization must be performed. The frontpanel value $\mathcal{T}_\mathrm{f}[i, j]$ is replaced with $\log(T_\mathrm{max})$, if the trial solution exceeded the upper bound, or with $\log(T_\mathrm{min})$, if it exceeded the lower bound, and the function is minimized again with respect to $\mathcal{T}_\mathrm{b}[i, j]$. From the first equation of (6.36), we obtain

$$\mathcal{T}_\mathrm{b}[i, j] \leftarrow \frac{k_\mathrm{p} S[i, j] + h^2\, w_\mathrm{ad}[i, j]\big(\log(L_\mathrm{in}[i, j]/B) - \mathcal{T}_\mathrm{f}[i, j]\big)}{4k_\mathrm{p} + h^2\, w_\mathrm{ad}[i, j]} \tag{6.38}$$

It can be noticed that the expressions (6.37) and (6.38) do not depend on the previous value of $\mathcal{T}_\mathrm{f}[i, j]$; therefore, it is not necessary to store the frontpanel and the memory requirements are reduced. Modern processors include instructions which compute the maximum or minimum of two numbers without using a branch which may reset the pipeline and introduce a performance penalty. Depending on the hardware, it may be convenient to replace the test on the frontpanel with max/min instructions and compute step (6.38) in any case; this last operation will simply leave $\mathcal{T}_\mathrm{b}$ unchanged if no constraint is active. Finally, it can be noticed that the denominator in (6.38) remains constant in subsequent iterations; therefore, it may be convenient to precompute its reciprocal, thus replacing the division with a simpler multiplication, provided that the additional memory requirements and latency are acceptable.

As shown in the previous Subsections, the proposed algorithm is parallelized by using the red-black ordering and accelerated using the multigrid technique. The input image $\log(L_\mathrm{in}/B)$

and the weights w_{ad} are downsampled by a factor 2 using a Galerkin projection and the low-resolution backpanel is computed (the frontpanel can be ignored as previously noticed). Since the constraints on the frontpanel are constant over the whole image, they are left unaltered. The backpanel is then interpolated back to the original resolution and a few iterations of the modified red-black Gauss-Seidel method are performed. We have verified that $5 \div 10$ iterations are enough for convergence. The downsampled problem is processed recursively until a sufficiently small image is obtained. In our implementation we stop the recursion when the image is less than 4 pixels wide or tall, because a smaller size would require a different treatment of the boundary conditions without introducing a noticeable improvement in the convergence speed. At the lowest level, the backpanel is simply initialized with the square root of the input image. It is not advisable to solve this lowest-level equation exactly using a quadratic programming solver because, since the input image has been smoothed out significantly by the repeated downsampling, the quadratic programming problem is likely to admit multiple solutions which can give rise to a numerical instability.

We conclude this Section with the observation that, as the parameter k_{p} is reduced, the reconstruction error E_{rec} also decreases because a lower weight is assigned to the competing objective E_{par}. Intuitively, if we take the limit $k_{\mathrm{p}} \to 0$, perfect reconstruction is achieved and the solution coincides with the one produced by the Method 2 described in Section 6.6. The expression (6.38) used in the proposed numerical algorithm is well defined also for $k_{\mathrm{p}} = 0$; therefore it is possible to use the same code also for Method 2, at the cost of a few redundant operations, and both methods can be considered as belonging to the same framework.

6.10 Experimental results

We designed a test program that splits an image using the proposed method and displays it onto a prototype dual layer LCD display. The user can change the algorithm parameters interactively and perform a visual assessment of the results. In this test implementation, we computed the adaptation level L_{ad} in (6.18) by means of a linear IIR lowpass filter. The test program is written in the C programming language, with some critical portions optimized in Assembly, and the splitting of a 512×512 pixel image takes approximately 0.14 seconds on a 1.5 GHz Pentium 4 processor. Processing at standard video rates is within reach if a suitable hardware is used; however it should be mentioned that the present applications of the proposed technique are primarily intended for the field of static image display.

Our preliminary tests showed that, at least with the selected medical images, parameter variations produced a visible change in the backpanel and frontpanel image when viewed separately, but no perceivable change in the complete dual-layer image. This shows that the reconstruction error is below the threshold of visibility. If the TVI function is included in E_{rec}, a larger reconstruction error is allowed in the dark portions of the image, since the model predicts a lower sensitivity of the human eye in this case. Consequently, a smoother backpanel image was observed in those points, thus confirming that a relaxation of the perfect reconstruction constraint provides a greater number of degrees of freedom. Also in this case, the difference was visible only when viewing the two panels separately. A sample original and processed image is shown in Figure 6.14. It should be noticed that the original (leftmost image) is a 16-bit image, which can not be easily reproduced on paper or on a conventional display; here it has been mapped in order to reduce its dynamic range.

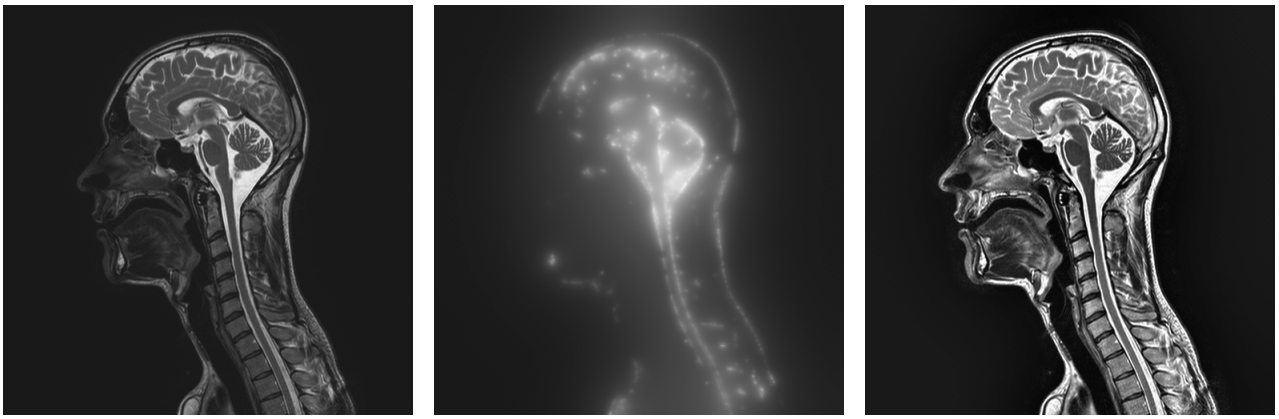

Figure 6.14: Example of the proposed splitting algorithm. Left to right: source image, backpanel, frontpanel. The splitting was performed using $k_{\mathrm{p}} = 0.1$ in (6.14) and Ferwerda's TVI function in (6.18). The source image has been mapped in order to reduce its dynamic range.

On the tested medical images, our method appears to remove completely the parallax error and no artifacts were visible when the display was observed off-axis. On some synthetic test images containing thin white lines over a uniform dark background, some halos were still visible near the edge. On the other hand, artifacts are intrinsically more visible on synthetic images because no masking occurs; the edges in actual images are generally not as sharp, and are often surrounded by textures that mask the visibility of halos.

In order to give a quantitative measure of the reconstruction error, we computed the expression

$$E_{\mathrm{ad}}(x,y) \triangleq \frac{B\,T_{\mathrm{b}}(x,y)\,T_{\mathrm{f}}(x,y) - L_{\mathrm{in}}(x,y)}{\mathrm{TVI}\big(L_{\mathrm{ad}}(x,y)\big)}, \tag{6.39}$$

which was used in Section 6.7 to define the functional E_{rec}. The physical dimension of $E_{\mathrm{ad}}(x,y)$ is just noticeable differences (JNDs), since the expression (6.39) contains the ratio of the luminance difference and the TVI function. The output of the splitting algorithm, and a false-color representation of the reconstruction error computed using (6.39), are shown Figure 6.15. We used $k_{\mathrm{p}} = 0.1$ and Barten's TVI function calibrated for a maximum luminance (white level) of 500 cd/m^2. For this test image, the reconstruction error introduced using the proposed method has an average value of approximately 0.3 JNDs. The parallax error can be estimated by computing the normalized gradient $\|\nabla T_{\mathrm{b}}/T_{\mathrm{b}}\|$ used in the definition of E_{par} in Equation (6.15) or (6.16). With the same choice of the parameters used in Figure 6.15, the normalized gradient takes a mean value of 0.044.

In Table 6.2 we have reported the parallax and reconstruction errors obtained with the proposed method for different choices of the parameter w, which sets the tradeoff between the parallax and reconstruction errors in the proposed variational splitting algorithm. The other parameters in the algorithm take the same values used in the previous examples. As expected, an increase in w privileges a faithful reconstruction at the cost of a steeper backpanel.

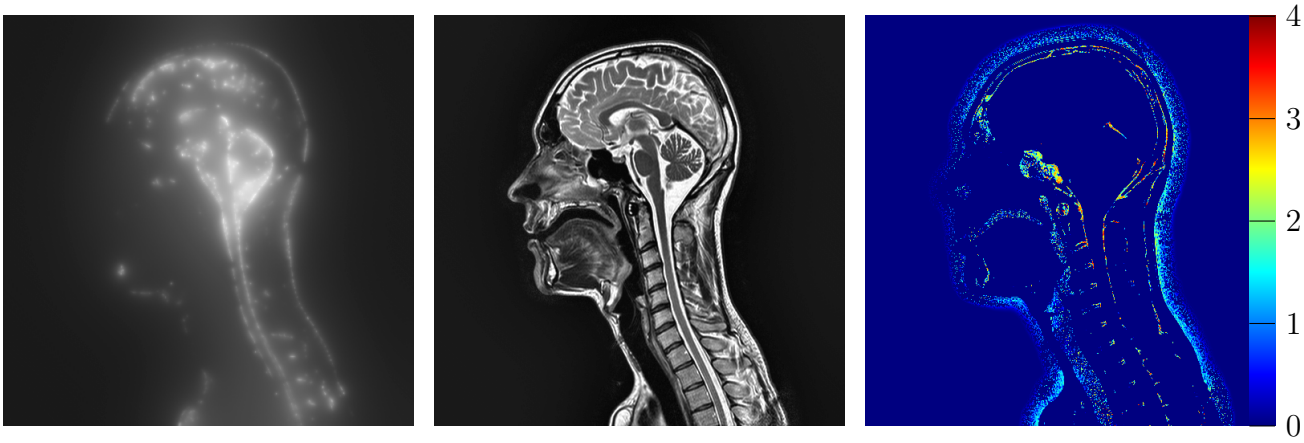

Figure 6.15: Example of backpanel, frontpanel and reconstruction error (measured in JNDs) obtained using the proposed method.

k_p	Mean $\|\nabla T_\mathrm{b}/T_\mathrm{b}\|$	Mean E_ad
10^{-1}	0.036	1.55
10^{-2}	0.039	0.67
10^{-3}	0.044	0.34
10^{-4}	0.051	0.18

Table 6.2: Parallax error (adimensional) and reconstruction error (JNDs) for different values of the weight factor w

6.11 Minimization of in-panel glare

Some measurements, conducted independently in the laboratories of the Center for Devices and Radiological Health, US Food and Drug Administration, under the supervision of Dr. Aldo Badano, showed that the unwanted scattering of light in the space between the two panels can cause in some cases a failure of the multiplicative law (6.1). More precisely, if the backpanel is bright and the frontpanel is dark in the same location, the light blocked by the frontpanel is scattered sideways inside the sandwich structure and can increase the brightness of nearby pixels. This artifact, which we shall call *in-panel glare*, was only measured with a collimated radiometer probe on a specifically designed synthetic test image and was not visible by naked eye on the medical images we used so far in our experiments; nevertheless we chose to introduce, for testing purposes, a modification in the splitting algorithm which addresses this issue. Since the described glare artifacts occurs in the points where the light is transmitted by the backpanel and blocked by the frontpanel, we added a third objective to the functional (6.23) which privileges a dark backpanel image in order to block the light as soon as possible and prevent unwanted paths. Computational considerations suggest to use the integral of the logarithm of the backpanel transmittance. The modified functional becomes therefore

$$\iint \left\{ \frac{k_\mathrm{p}}{2} \left\| \nabla \mathcal{T}_\mathrm{b} \right\|^2 + \frac{w_\mathrm{ad}}{2} \left[\mathcal{T}_\mathrm{b} + \mathcal{T}_\mathrm{f} - \log(L_\mathrm{in}/B) \right]^2 + k_\mathrm{g} \mathcal{T}_\mathrm{b} \right\} dx\, dy, \tag{6.40}$$

where k_g is an user-adjustable parameter which sets the relative weight of this third objective. The parameter k_g can also take negative values; in this case the splitting algorithm will privilege a brighter backpanel.

Unfortunately, the proof of Section 6.8 can not be generalized to the modified functional (6.40), and indeed we have verified that the backpanel can hit a constraint if $k_g \neq 0$. Therefore, the numerical method described in Section 6.9 needs some modifications. Following the same general principles, we discretize the functional and perform a point-by-point minimization, using a multigrid method to generate a suitable starting guess. For each pixel, we first compute a trial solution by performing an unconstrained minimization of the discretized functional with respect to $\mathcal{T}_b[i,j]$ and $\mathcal{T}_f[i,j]$. This requires the solution of the following linear system, which is similar to the one of Equation (6.36) with the addition of the k_g term:

$$\begin{cases} h^2\, w_{ad}[i,j]\big(\mathcal{T}_b[i,j] + \mathcal{T}_f[i,j] - \log(L_{in}[i,j]/B)\big) + h^2 k_g = k_p\big(S[i,j] - 4\mathcal{T}_b[i,j]\big) \\ h^2\, w_{ad}[i,j]\big(\mathcal{T}_b[i,j] + \mathcal{T}_f[i,j] - \log(L_{in}[i,j]/B)\big) = 0 \end{cases} \tag{6.41}$$

The trial solution we obtain is

$$\mathcal{T}_b[i,j] \leftarrow \frac{1}{4}\left(S[i,j] - \frac{h^2 k_g}{k_p}\right) \qquad \mathcal{T}_f[i,j] \leftarrow \log(L_{in}[i,j]/B) - \mathcal{T}_b[i,j] \tag{6.42}$$

If both $\mathcal{T}_b[i,j]$ and $\mathcal{T}_f[i,j]$ satisfy the bounds, the trial solution (6.42) is accepted and the next pixel is processed; otherwise, a constrained minimization is performed.

The first method we tried computes the backpanel value $\mathcal{T}_b[i,j]$ using the left expression in (6.42), clips it if the value falls outside the bounds and then computes the frontpanel value $\mathcal{T}_f[i,j]$ using the right expression in (6.42). If the frontpanel value satisfies the constraints, it is accepted and the next pixel is processed. If instead $\mathcal{T}_f[i,j]$ falls outside the bounds, the algorithm tries to set $\mathcal{T}_f[i,j]$ equal to the two extremal values $\log(T_{min})$ and $\log(T_{max})$ and minimizes the function again with respect to $\mathcal{T}_b[i,j]$. From the first equation of (6.41), we obtain

$$\mathcal{T}_b[i,j] \leftarrow \frac{k_p S[i,j] + h^2\, w_{ad}[i,j]\big(\log(L_{in}[i,j]/B) - \mathcal{T}_f[i,j]\big) - h^2 k_g}{4k_p + h^2\, w_{ad}[i,j]} \tag{6.43}$$

The backpanel is clipped again if this new value falls outside the bounds. Finally, the algorithm selects the correct solution by computing the Lagrange multipliers and verifying if their sign satisfies the conditions of the Karush-Kuhn-Tucker theorem. Unfortunately, our tests showed that this method has a poor numerical stability and rounding errors can cause the test on the Lagrange multipliers to fail for both solutions.

The second method we tried clips the frontpanel value $\mathcal{T}_f[i,j]$ if it falls outside the bounds and updates the backpanel value $\mathcal{T}_b[i,j]$ using the expression (6.43). Finally, the backpanel is examined and clipped if necessary. This second method does not suffer from the aforementioned numerical instability.

6.12 Conclusions and future work

In this Chapter we presented a prototype dual layer LCD display capable of surpassing the dynamic range of the conventional radiographic film. The display uses two liquid crystal panels stacked one on top of the other in order to modulate the light coming from the backlight unit twice and obtain a significantly darker black level. In order to operate the display, an appropriate splitting algorithm is required for the generation of the two images which drive the backpanel and the frontpanel. The splitting algorithm should minimize the parallax error

introduced by the viewing angle, the visibility of reconstruction error, and optionally the in-panel glare caused by the unwanted scattering of light between the two panels. Appropriate constraints are also necessary in order to guarantee that the images generated by the splitting algorithm fall between the black and white levels of the panels. We proposed different algorithms which successively incorporate the different objectives. The methods we proposed are based on constrained optimization techniques, which allow a great flexibility in the design of the objectives and constraints; we then proposed an efficient numerical algorithm, based on multigrid methods, which can solve this apparently complex problem at a remarkably low computational cost. The algorithm is able to satisfactorily reduce the parallax error on medical images also in presence of sharp edges. Perfect reconstruction is achieved when possible, and is lost only in the points where the limited dynamic range of the panels would cause the frontpanel image to saturate. In this case, the visibility of the reconstruction error is minimized by means of a specific-designed visible difference metric, which is at the same time simple enough to allow an efficient computation and accurate enough to produce high-quality results.

The prototypes and the image processing algorithms are still under development. Concerning the hardware aspect, it might be possible to increase the total brightness by using a backpanel with a lower resolution, which typically has a higher transmittance because a smaller fraction of area is occupied by the electrical wirings. This is not possible yet with the current prototypes because the 1.3-megapixel (1280x1024 pixel) panels we are using have a 5:4 aspect ratio which is not available at lower resolutions. However, we are planning to build 5-megapixel prototypes, which also have a 5:4 aspect ratio, and in this case a 5-megapixel frontpanel can be coupled with a 1.3-megapixel backpanel. It should be noticed that this technique only aims at increasing the brightness, and is not likely to reduce the computational cost of the splitting algorithm, because the frontpanel must be sharpened in any case taking care to prevent the clipping. For the same reason, it is not advisable to introduce a diffuser foil between the two layers in order to blur the backpanel. On the contrary, a diffuser would absorb some light, thus reducing the total brightness, and further increase the distance between the panels and the consequent parallax error. Concerning the software aspect, one topic which could be investigated is the replacement of the red-black Gauss-Seidel method with different iterative solvers, such as for instance the Conjugate Gradient (CG) method and the Alternating Directions Implicit (ADI) method [57]. These methods typically require a smaller number of iterations, however each iteration has a significantly higher complexity, and the extension of these methods to constrained problems is not straightforward. A multigrid method will still be used to generate a coarse approximtion of the solution and improve the convergence.

Chapter 7

Equivalent transfer curve and bit depth of a dual-layer LCD display

7.1 Outline

A conventional digital display is able to reproduce a discrete number of luminance levels. This number – which depends on the bit depth of the input signal and on the hardware capabilities – is well defined and gives a measure of the grayscale reproduction accuracy of the device. In a dual layer display [76, 24], instead, the luminance levels obtained by the combination of frontpanel and backpanel transmittances are partially overlapped and nonuniformly spaced; therefore it is difficult to define and compute a figure which represents the "bit depth" of the device.

In this Chapter, we attempt to estimate the equivalent number of levels of a dual layer display by means of a statistical analysis of the quantization noise. We derive an expression which gives the noise variance as a function of the digital driving levels (DDLs) of the frontpanel and backpanel. A constrained optimization of this function yields the DDL pair which produces a desired luminance output with the minimum noise variance. Finally, a numerical integration gives a statistical estimate of the number of equivalent gray levels the display is able to reproduce.

7.2 Noise model of a digital display

A digital display is able to reproduce a discrete set of luminance levels. If q is the digital driving level (DDL) – which for an 8-bit display can take integer values between 0 and 255 – the output luminance is given by $L_{\text{out}} = T(q)$, where $T(\cdot)$ is a nonlinear transfer function commonly known as *gamma curve*. The limited set of displayable values imposes some quantization. If a simple rounding or truncation is used for this purpose, the output is a deterministic function of the input and contains therefore a systematic error or bias. Moreover, a statistical analysis of the quantization error is difficult in this case due to the strong nonlinearity introduced by the truncation or rounding operation.

In order to avoid these problems, some random component must be introduced, and a commonly used technique is *dithering*. Rather than limiting q to integer values, we allow it to take any

real value in the range $[0, 255]$. This number can be split into its integer and fractional parts:

$$q = q_{\text{int}} + q_{\text{frac}} \qquad q_{\text{int}} \in \mathbb{Z},\ 0 \leq q_{\text{frac}} < 1 \tag{7.1}$$

We then set the output to q_{int} with probability $1 - q_{\text{frac}}$ or to $q_{\text{int}} + 1$ with probability q_{frac}. This operation is typically implemented in practice by adding to q a white *dithering noise* n_D with a uniform probability density function (PDF) between $-1/2$ and $1/2$ and rounding the result to the nearest integer; some computations can show that the two definitions are equivalent. If dithering is used, the quantized output becomes a random variable; some computations can show that its mean value is equal to the input q (and consequently no bias is present on average), and that the quantization noise is white and not correlated to the input. It should be noticed, however, that the quantization noise is not statistically independent from the input, because the output can only take integer values[1]. If we assume that the fractional part q_{frac} has a uniform distribution between 0 and 1 and compute an average, we obtain that the quantization noise has a triangular PDF between -1 and 1. Therefore, the effect of dithering can be *approximated* – if only its PDF and correlation are considered – by an additive white noise n with a triangular PDF that we shall indicate with $f(n)$. It is easy to show that this kind of noise has mean value $E\{n\} = 0$ and variance $\sigma^2 \triangleq E\{n^2\} = 1/6$. A block diagram of the system, which includes the dithering and the nonlinear mapping $T(\cdot)$, is shown in Figure 7.1.

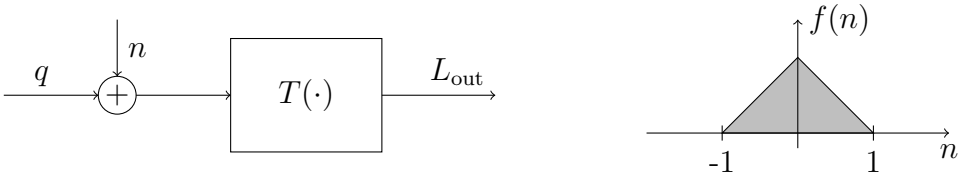

Figure 7.1: Noise model of a single-layer display.

The output luminance of a dithered display is a random variable with the expression $L_{\text{out}} = T(q + n)$. In order to calculate its mean and variance, we shall approximate the transfer curve $T(\cdot)$ with its first-order Taylor expansion around q, i.e. $L_{\text{out}} \approx T(q) + T'(q)\, n$. The mean value of L_{out} is

$$E\{L_{\text{out}}\} = \int \left[T(q) + T'(q)\, n\right] f(n)\, dn = T(q). \tag{7.2}$$

In order to reproduce a desired luminance L_{in} without bias, we request that $E\{L_{\text{out}}\} = L_{\text{in}}$. This is accomplished by taking $q = T^{-1}(L_{\text{in}})$ and quantizing it with dithering. The variance of the output luminance is

$$E\{(L_{\text{out}} - T(q))^2\} = \int \left[T(q) + T'(q)\, n - T(q)\right]^2 f(n)\, dn = T'^2(q)\, \sigma^2. \tag{7.3}$$

Due to the nonlinear mapping $T(\cdot)$, the variance of L_{out} is not constant but depends on the input.

Following a similar reasoning, a dual-layer display can be modeled as shown in Figure 7.2. The DDL values of the frontpanel and backpanel are q_F and q_B respectively, and the

[1]The quantization noise can be made statistically independent from the input with the technique of *subtractive dithering*, in which the *same* noise process n_D is added to the signal q before the rounding and subtracted after the rounding. However, this technique is inapplicable in most practical cases.

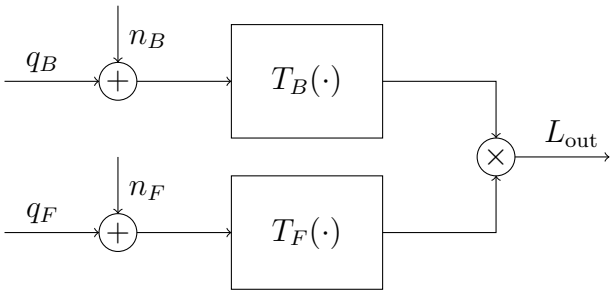

Figure 7.2: Noise model of a dual-layer display.

output luminance is a random variable with expression $L_{\text{out}} = T_F(q_F + n_F)\, T_B(q_B + n_B) \approx \big[T_F(q_F) + T'_F(q_F)\, n_F\big]\big[T_B(q_B) + T'_B(q_B)\, n_B)\big]$. The transfer curves $T_F(\cdot)$ and $T_B(\cdot)$, of the frontpanel and backpanel respectively, can be different in general. The quantization noises n_F and n_B are statistically independent and have the same triangular PDF for both panels, because in both cases a real number is being quantized to 8 bits with dithering. However, the backpanel appears slightly blurred due to the presence of the frontpanel, and this has the effect of smoothing out the noise and reducing its variance. Therefore, we shall indicate the variances with two different symbols σ_F^2 and σ_B^2. The mean value of L_{out} is

$$\begin{aligned} E\{L_{\text{out}}\} &= \iint \big[T_F(q_F) + T'_F(q_F)\, n_F\big]\big[T_B(q_B) + T'_B(q_B)\, n_B)\big]\, f(n_F)\, f(n_B)\, dn_F\, dn_B = \\ &= T_F(q_F)\, T_B(q_B). \end{aligned} \tag{7.4}$$

For simplicity of notation, we shall drop (q_F) and (q_B) in the following. The variance of L_{out} is

$$\begin{aligned} E\{(L_{\text{out}} - T_F T_B)^2\} &= \iint \big[(T_F + T'_F n_F)(T_B + T'_B n_B) - T_F T_B\big]^2\, f(n_F)\, f(n_B)\, dn_F\, dn_B = \\ &= (T_F^2 + T_F'^2 \sigma_F^2)(T_B^2 + T_B'^2 \sigma_B^2) - T_F^2 T_B^2 \end{aligned} \tag{7.5}$$

7.3 Optimal splitting in absence of parallax error

If no parallax error is considered, the splitting can be performed on a pixel-by-pixel basis. In order to reproduce a desired luminance L_{in} without bias, the condition $T_F T_B = L_{\text{in}}$ should be met. We seek the pair (T_F, T_B) which minimizes the expression (7.5) subject to the constraint $T_F T_B = L_{\text{in}}$. The constraint is added by introducing a Lagrange multiplier λ:

$$\mathcal{L} \triangleq (T_F^2 + T_F'^2 \sigma_F^2)(T_B^2 + T_B'^2 \sigma_B^2) - T_F^2 T_B^2 + \lambda(T_F T_B - L_{\text{in}}) = \text{Min} \tag{7.6}$$

The expression (7.6) contains the nonlinear functions $T_F(\cdot)$ and $T_B(\cdot)$, therefore some iterative method is needed. Newton's method is difficult to use in this case, because it requires the knowledge of the second derivatives of the transfer curves. As a workaround, we shall replace $T_F(\cdot)$ and $T_B(\cdot)$ in the expression (7.6) with their first-order Taylor expansion, or equivalently, we shall minimize (7.6) with respect to T_F and T_B assuming that T'_F and T'_B are constant. By taking the partial derivatives of (7.6) with respect to T_F, T_B and λ and setting them equal to

zero, we obtain the following nonlinear system:

$$\begin{cases} \partial\mathcal{L}/\partial T_F = 2(\sigma_B T'_B)^2 T_F + \lambda T_B = 0 \\ \partial\mathcal{L}/\partial T_B = 2(\sigma_F T'_F)^2 T_B + \lambda T_F = 0 \\ \partial\mathcal{L}/\partial\lambda \;\; = T_F T_B - L_{\text{in}} = 0 \end{cases} \tag{7.7}$$

The system has four solutions; the only real and positive one is

$$T_F = \sqrt{L_{\text{in}} \frac{\sigma_F T'_F}{\sigma_B T'_B}} \qquad T_B = \sqrt{L_{\text{in}} \frac{\sigma_B T'_B}{\sigma_F T'_F}} \tag{7.8}$$

The values of q_F and q_B are obtained by inversion of the transfer curves, and the new values of T'_F and T'_B are computed. This procedure is then iterated until convergence. If either T_F or T_B falls outside the dynamic range of the panel, it is saturated and the other variable is computed by division.

In Figure 7.3 we show an example of the optimal splitting curves computed using the proposed method. In the left plot, we used for both T_F and T_B the curve defined in the international standard recommendation ITU-R BT.709, which has an analytical expression and allows us to compute the derivative easily. This curve has a dynamic range of approximately 200, therefore the dual-layer display could theoretically reach a dynamic range of 40000. In the right plot, we used DICOM-calibrated panels with a dynamic range of 500 and a maximum luminance of 500 cd/m^2; such a display could theoretically reach a dynamic range of 250000.

If we use the same values for σ_F and σ_B, the solution is $T_F = T_B = \sqrt{L_{\text{in}}}$ by symmetry and the plot is a line with slope 0.5 in log-log scale. If instead we use a lower value for σ_B, the optimal splitting algorithm assigns different values to the frontpanel and the backpanel; moreover, for dark and bright values of L_{in}, one panel saturates and the image is entirely displayed on the other one. In the computations presented in this Chapter, we arbitrarily set $\sigma_B = \sigma_F/2$; the actual value should be measured experimentally.

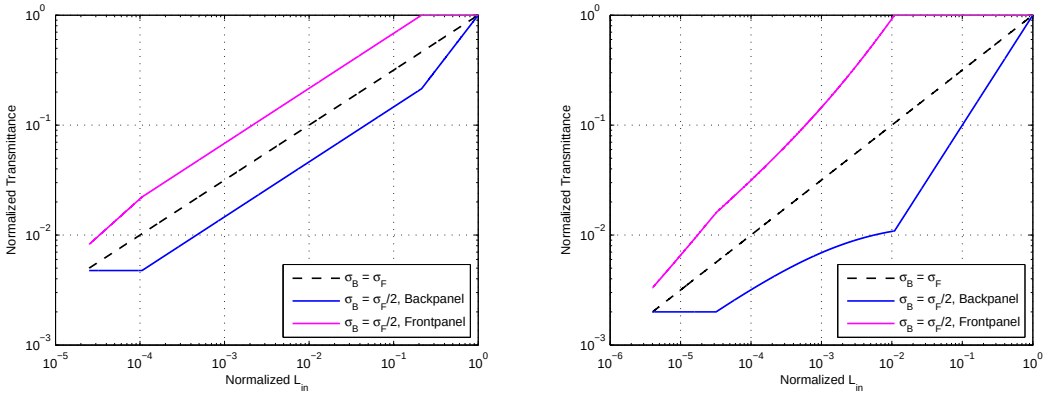

Figure 7.3: Examples of optimal splitting with ITU-R BT.709 (left) and DICOM (right) transfer curves. Each plot shows the results obtained with $\sigma_F = \sigma_B$ (which produces a symmetric splitting $T_F = T_B = \sqrt{L_{\text{in}}}$) and with $\sigma_B = \sigma_F/2$ (which produces an asymmetric output).

7.4 Equivalent transfer curve

In order to compute the equivalent bit depth of a dual-layer display, we generate an equivalent transfer curve $T_{\mathrm{eq}}(q)$ of a theoretical single-layer display which produces, for every luminance level L_{in}, the same quantization noise that the dual-layer display would produce when the optimal splitting described in Section 7.3 is used. Following the expression derived in (7.3), we request that the equivalent transfer curve satisfy the differential equation

$$\sigma_{\mathrm{opt}}(T_{\mathrm{eq}}(q)) = T'_{\mathrm{eq}}(q)\,\sigma, \tag{7.9}$$

where $\sigma_{\mathrm{opt}}(\cdot)$ is the standard deviation of the quantization noise as a function of luminance obtained using the optimal splitting and $\sigma = \sqrt{1/6}$ is the standard deviation of the triangular dithering noise according to the model in Figure 7.1. In the points where no panel saturates, σ_{opt} is computed using (7.5). In the points where one panel saturates, instead, we shall assume that it does not introduce quantization noise, because the panel displays a fixed DDL of 0 or 255 without dithering; therefore only the other (unconstrained) panel contributes to σ_{opt}. As initial condition we set $T_{\mathrm{eq}}(0) = L_{\mathrm{min}}$, where L_{min} is the black level of the dual-layer display. The curve $T_{\mathrm{eq}}(q)$, obtained by a numerical integration of (7.9), will reach the white level L_{max} in correspondence of a DDL value q_{max}, which is expectedly greater than 255 because the range of the curve is larger and the quantization noise is lower compared to a single-layer display. This value q_{max} such that $T_{\mathrm{eq}}(q_{\mathrm{max}}) = L_{\mathrm{max}}$ represents the number of equivalent levels of the dual-layer display following the model described in Section 7.2. It should be noticed that this number represents an upper limit, because the optimal splitting is being used; real-life applications will use a splitting algorithm which reduces the parallax error, but increases the quantization noise due to the suboptimal choice of the pair (T_F, T_B).

The equivalent transfer curves corresponding to the optimal splittings of Figure 7.3 are shown in Figure 7.4. The left plot shows the results with the ITU-R BT.709 transfer curve and a total dynamic range $L_{\mathrm{max}}/L_{\mathrm{min}} = 40000$. With these settings, the number of equivalent levels is approximately 360 with $\sigma_F = \sigma_B$ and 625 with $\sigma_B = \sigma_F/2$. A lower value of σ_B reduces the quantization noise significantly, therefore an accurate measurement of this parameter is necessary in order to have reliable results. The right plot shows the results with the DICOM transfer curve and a total dynamic range of 250000. Despite the higher dynamic range of the display, the number of equivalent gray levels remains approximately unchanged: 360 with $\sigma_F = \sigma_B$ and 675 with $\sigma_B = \sigma_F/2$.

In order to better visalize the distribution of the equivalent levels, it is possible to draw a histogram. Figure 7.5, for instance, shows the number of equivalent levels per decade of luminance obtained in the 4 examples above.

7.5 Visibility of the quantization noise

The computations in the previous sections only take into account the physical nature of the display, and do not consider the human perception of the quantization noise. In order to measure the actual visibility of the noise at different luminance levels L, we divide the luminance difference between two adjacent levels by the just noticeable difference (JND). The luminance difference is approximately equal to the derivative $T'_{\mathrm{eq}}(q)$ of the equivalent transfer curve, scaled

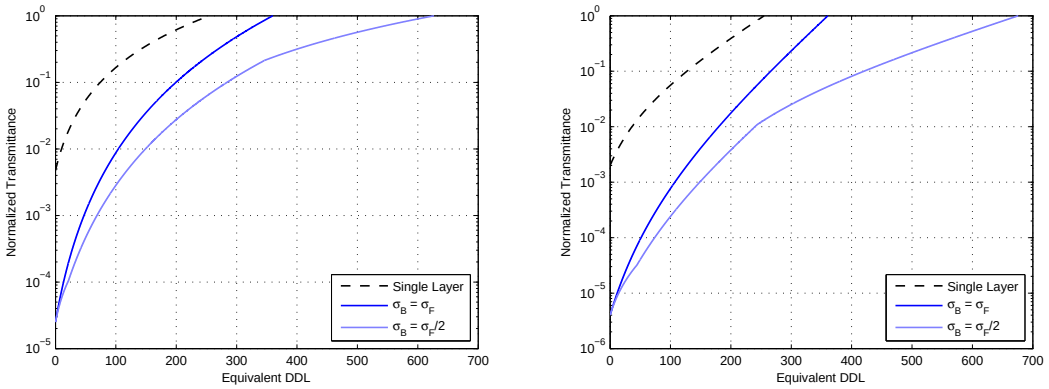

Figure 7.4: Equivalent transfer curves of dual-layer displays with ITU-R BT.709 (left) and DICOM (right) panels.

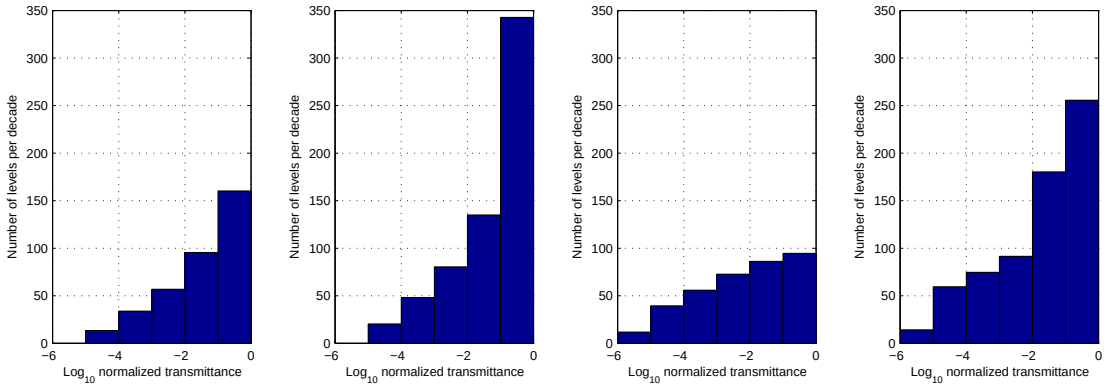

Figure 7.5: Histograms of the number of levels per decade of luminance obtained with different optimal splittings. Left to right: ITU-R BT.709 panels with $\sigma_B = \sigma_F$ and $\sigma_B = \sigma_F/2$, DICOM panels with $\sigma_B = \sigma_F$ and $\sigma_B = \sigma_F/2$.

in physical luminance units; the JND is computed using a threshold-versus-intensity (TVI) function. In Figure 7.6 we plotted the results of this computation with the same settings as the previous examples, using Barten's model for the TVI function with a white level of 500 cd/m^2. As a comparison, we also plotted the quantization noise of a single-layer display computed with (7.3).

With ITU-R BT.709 panels (left plot), the quantization noise of the dual-layer display decreases in the dark and bright areas. For the single-layer display, instead, the noise is significantly higher and increases in the dark areas. It can be noticed that, with $\sigma_F = \sigma_B$, the dual-layer display has a slightly worse behavior than the single-layer for luminances close to the white level. This is possibly due to the fact that both panels operate near the white level and the benefit of the backpanel is outdone by the additional quantization noise it introduces. With $\sigma_B = \sigma_F/2$, instead, the quantization noise drops for dark and bright luminance levels, due to the assumption we made in Section 7.4 that no dithering is performed on the panels which are saturated to their black or white level.

With DICOM panels (right plot), the visibility of the quantization noise in the single-layer is

constant by construction. Surprisingly, with $\sigma_F = \sigma_B$, the dual-layer behaves worse than the single-layer in a significant range. This is possibly due to the fact that, in the bright areas, the DICOM curve behaves worse than the ITU-R BT.709 curve as shown in the plots. With $\sigma_B = \sigma_F/2$, instead, the noise is significantly lower.

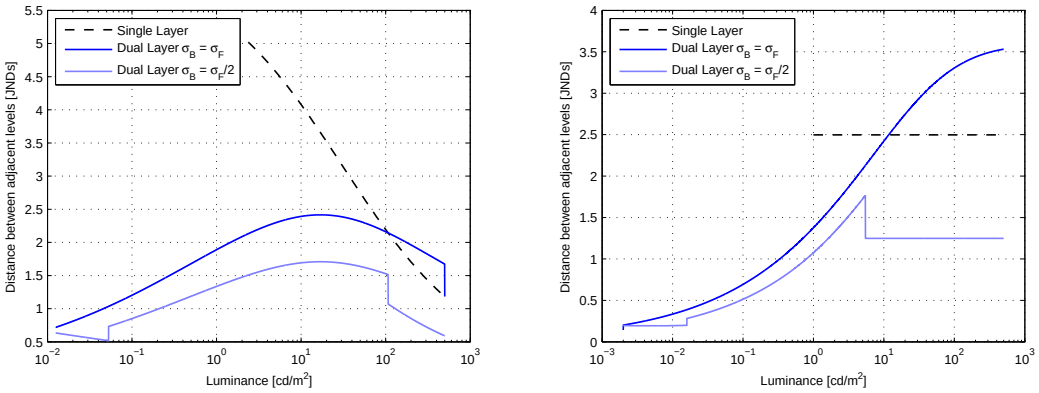

Figure 7.6: Visibility of the quantization noise in dual-layer displays with ITU-R BT.709 (left) and DICOM (right) panels.

7.6 Influence of the panel transfer curve

It is reasonable to investigate how the number of equivalent levels, computed with the method proposed in Section 7.4, depends on the transfer curves T_F and T_B of the panels. For the moment, we shall assume that the two panels have equal transfer curves and noise variances. This simplifies the computation because the optimal splitting is a square root; the computation for the general case $\sigma_F \neq \sigma_B$ has not been performed yet.

For equal panels and square-root splitting, the variance of the quantization noise at a given luminance level L, according to the expression (7.5), reduces to

$$\sigma^2_{\text{opt}}(L) = 2\sigma^2\, T^2(q)\, T'^2(q) + \sigma^4\, T'^4(q) \qquad \text{with } T(q) = \sqrt{L} \tag{7.10}$$

By integrating (7.9), we obtain that the number of equivalent levels is

$$q_{\max} = \int_{L_{\min}}^{L_{\max}} \frac{\sigma\, dL}{\sigma_{\text{opt}}(L)} \tag{7.11}$$

If we make the change of variables $L = T^2(q)$, where q is the DDL value of the individual panels (which is equal to $T^{-1}(\sqrt{L})$ by hypothesis), and replace σ_{opt} with its expression (7.10), we obtain

$$q_{\max} = \int_0^{255} \frac{2\sigma\, T(q)\, T'(q)\, dq}{\sqrt{2\sigma^2\, T^2(q)\, T'^2(q) + \sigma^4\, T'^4(q)}} = \int_0^{255} \sqrt{2} \left[1 + \frac{\sigma^2\, T'^2(q)}{2T^2(q)}\right]^{-1/2} dq \tag{7.12}$$

The term in square brackets is always greater than one, therefore we obtain the inequality

$$q_{\max} < \int_0^{255} \sqrt{2}\, dq = 255\sqrt{2} \approx 361 \tag{7.13}$$

In order to compute a lower limit, we can observe that $(1+x)^{-1/2} \geq 1 - x/2$, therefore

$$q_{\max} > \int_0^{255} \sqrt{2} \left[1 - \frac{\sigma^2 \, T'^2(q)}{4T^2(q)} \right] dq \tag{7.14}$$

By using techniques of Calculus of Variations, it is possible to show (the computation is omitted here) that the functional (7.14) – with boundary conditions $T(0) = 1/d$ and $T(255) = 1$, where d is the dynamic range of the individual panels – is minimized by the function

$$T(q) = \frac{1}{d} \exp\left(\frac{q \log d}{255} \right). \tag{7.15}$$

By substituting (7.15) into (7.14) and computing the integral, we obtain the lower bound

$$q_{\max} > \int_0^{255} \sqrt{2} \left[1 - \left(\frac{\sigma \log d}{2 \cdot 255} \right)^2 \right] dq = 255\sqrt{2} \left[1 - \left(\frac{\sigma \log d}{2 \cdot 255} \right)^2 \right] \tag{7.16}$$

For $d = 500$, for instance, the term in square brackets is equal to $1 - 2.47 \cdot 10^{-5}$, therefore the lower and upper bound almost coincide. We conclude therefore that, with equal panels and square-root splitting, the equivalent number of levels is, to an excellent approximation, independent of the panels used in the dual layer display.

7.7 DICOM-compliant dual layer display

In a DICOM-compliant display, the luminance difference between two adjacent levels is a constant multiple of the JND. This requirement is not readily applicable to a dual layer display, because the output levels can not be defined unambiguously; it is however possible to request that the standard deviation of the quantization noise be a constant multiple of the JND: $\sigma_{\text{opt}}(L) = k \ \text{TVI}(L)$. We again make the assumption that the two panels are equal and a square-root splitting is used. By manipulating Equation (7.10), we obtain

$$\sigma_{\text{opt}}(L) = \sqrt{2\sigma^2 \, T^2(q) \, T'^2(q) + \sigma^4 \, T'^4(q)} = \sqrt{2}\sigma \, T(q) \, T'(q) \sqrt{1 + \frac{\sigma^2 \, T'^2(q)}{2T^2(q)}} \tag{7.17}$$

The second term under the square root is negligible, as shown in the previous Section, therefore the requirement of DICOM compliance becomes approximately

$$\sqrt{2}\sigma \, T(q) \, T'(q) = k \ \text{TVI}\big(L_{\max} T^2(q)\big), \tag{7.18}$$

where the scale factor $L_{\max}$ is necessary to convert from normalized to physical luminance values. Equation (7.18) can be re-written as

$$\sqrt{2}\sigma \, \frac{d}{dq} \left(\frac{1}{2} T^2(q) \right) = k \ \text{TVI}\big(L_{\max} T^2(q)\big) \tag{7.19}$$

Equation (7.19) shows that DICOM compliance is achieved when the slope of the *square* of the panel transmittances is a constant multiple of the JND. Therefore, it is easy to conclude that a DICOM-compliant dual layer display – with the meaning described at the beginning of this Section – should have the panel transmittances proportional to the square root of a DICOM grayscale display function, with luminance range equal to the one of the dual layer display.

7.8 Results and conclusions

In this Chapter, we proposed a mathematical model which allows to estimate the equivalent number of levels of a dual layer display by means of a statistical analysis of the quantization noise. This computation is difficult to perform by other means, because the levels obtained by multiplying all possible combinations of the panel transmittances are partially overlapped and nonuniformly spaced. Some assumptions were still necessary; in particular the quantization was modeled as an additive white noise – which is approximately true if dithering is used – and the splitting was performed on a pixel-by pixel basis neglecting the parallax error.

The number of levels estimated with this method is significantly below the theoretical maximum value of 2^{16}. If both panels have the same behavior, the quantization noise is minimized when the square root of the input image is displayed on both of them. In this case, the performance is almost independent of the response curve and the dynamic range of the individual panels. However, it is reasonable to assume that the quantization noise in the backpanel is less visible due to the optical blurring introduced by the frontpanel. With this assumption, the optimal splitting becomes asymmetric and the number of equivalent gray levels increases significantly; therefore, the optical blurring introduced by the frontpanel should be measured accurately in order to obtain a reliable estimate. The overall quantization noise of the dual-layer display depends on the luminance, even when DICOM-calibrated panels are used. It is possible to achieve a "DICOM-compliant" performance – in the sense that the standard deviation of the quantization noise at each luminance level is a constant multiple of the JND – by adjusting the panel response curves appropriately.

The results presented in this Chapter should be handled with care, because the model still has some limitations and the assumptions and approximations made in order to complete the computations might not be plausible in a real operating environment. In particular, Figure 7.6 shows that, for bright luminance levels, a dual-layer display with "optimal" square-root splitting introduces a greater error than the single-layer display. This suggests that it may be convenient to let one panel saturate in this case. This solution, however, was not produced by the constrained optimization algorithm described in Section 7.3 because the objective function (7.6) we used does not take into account the saturation; the assumption that a saturated panel does not introduce quantization noise was only used *a posteriori* in the computation of σ_{opt}. Moreover, the saturation of one panel would produce a discontinuous splitting curve, and possibly visual artifacts as a consequence.

Chapter 8

Quality assessment in medical images

8.1 Introduction and motivation

The quality assessment of medical images has significantly different objectives and methodologies compared to conventional photographic images. In the medical field, the quality of an image is not related to aesthetical properties; therefore, issues such as naturalness, fidelity, the observer's preference or the artist's intent are unimportant or totally absent. As a consequence, image enhancement algorithms such as the ones described in Chapter 3 are generally not used in the medical field because the space-varying mapping they perform, although aesthetically pleasing, can lead to the formation of spurious details in a medical image. For this reason, objective quality metrics are preferred over subjective ones.

Image fidelity metrics, based either on simple measures such as the mean squared error (MSE) and the peak signal to noise ratio (PSNR), or on more complex operators which take into account the mechanisms in the human visual system [15, 47], pose some issues because a reference image is not available in most cases. For instance, computed tomography or magnetic resonance images are generated synthetically by means of rendering algorithms which reconstruct three-dimensional sections of the patient's body based on the data collected by appropriate arrays of sensors. Moreover, as described in Chapter 5, radiographic images visualize as luminance some data which originally had a different physical dimension – such as for instance a tissue density – and a suitable nonlinear mapping must necessarily be performed. This mapping often plays a fundamental role in medical imaging systems and can significantly improve the visibility of the details in the image; however, gray level mappings can heavily penalize most image fidelity metrics.

Medical images are tools used to perform a diagnosis: for instance, a radiographic image should allow the radiologist to detect a specific target, such as a lesion, nodule or tumor, in a short time and with a high probability of success, avoiding both the missed detection of a present target (false negative) and the erroneous detection of artifacts or spurious details (false positive). Quality metrics should therefore be *task based*, meaning that the quality of an image processing or display technology is linked to the average probability of correct decision obtained by the physician when performing a specific diagnosis on images displayed using that technology. During the present research, we performed a set of observer studies, in cooperation with the Center for Devices and Radiological Health of the US Food and Drug

Administration, in order to evaluate and quantify the benefits introduced by the dual layer LCD displays introduced in Chapter 6. This technology is capable of reproducing images with a significantly wider luminance range and with an increased bit depth. The specific task we considered is the detection of a nodule or tumor on a mammographic image, because this kind of images, typically rich in details with very fine luminance differences, are likely to benefit from a high dynamic range display technology. The experimental method is described in detail in Section 8.2. In order to increase the number of test trials and reduce the experiment costs, it is possible in some cases to replace the human observer by appropriate software applications which simulate the diagnosis task by means of statistical classification techniques. In Section 8.3 we give a brief overview of this technique.

8.2 Experimental method

In a typical observer study, a sequence of test images is presented to a group of observers who are asked to perform a diagnosis, indicating whether they are able or not to detect a specific target in the image. The test images we used in our experiments are synthetic *clustered lumpy backgrounds* (CLBs) [10] which simulate the typical texture of a mammographic image. This technique allows to increase the number of trials significantly and provides noise-free images with an arbitrary bit depth; in this way, the test results are not influenced by possible quality defects in the source image and allow a more specific assessment of the display device and the image processing technique. The data set we used contains 532 images with a size of 128×128 pixels and a bit depth of 16 bits per pixel. The test target is a Gaussian-shaped signal, with a standard deviation of 10 pixels and relative density of 0.1, positioned in the middle of the image. Some example of images used in our test are visible in Figure 8.1. Since the target is added artificially to the image, a "ground truth" information is available and it is straightforward to verify if the observer's decision is correct or wrong. The density values of the image are then mapped to luminance using the technique described in Chapter 5 and displayed. The purpose of the study is to compare two presentation modes in order to determine which one provides the highest probability of correct diagnosis. In the present case, one mode corresponds to a high dynamic range image, displayed on a dual layer LCD prototype using the splitting algorithm described in Chapter 6; the other mode corresponds to a conventional display with a lower dynamic range and 8-bit quantization. The tests are performed using a dual layer prototype, and the conventional display is simulated by displaying the image on the frontpanel only over a white backpanel.

One possible inconvenience, which typically occurs in detection tasks such as the present one, is that the threshold of visibility for the test target does not have a well defined value. For signal amplitudes close to the threshold, the observer is generally not able to make a motivated decision and his answer is more likely the product of a random guess. For this reason, we designed a software application which performs the observer studies using the *two alternative forced choice* (2AFC) procedure [4, 10], which is specifically designed to address this issue. At each test trial, two different CLBs are selected randomly from the data set. On one of the two images, selected randomly, the target is superimposed. The two images are then processed with the algorithm under test and displayed simultaneously side by side to the user. The user must select the image which, in his opinion, contains the target, and his decision is recorded. The test trial is repeated multiple times until all the images in the data set are used; at the end of the test, the percentage of correct decisions is computed. No distinction is performed between

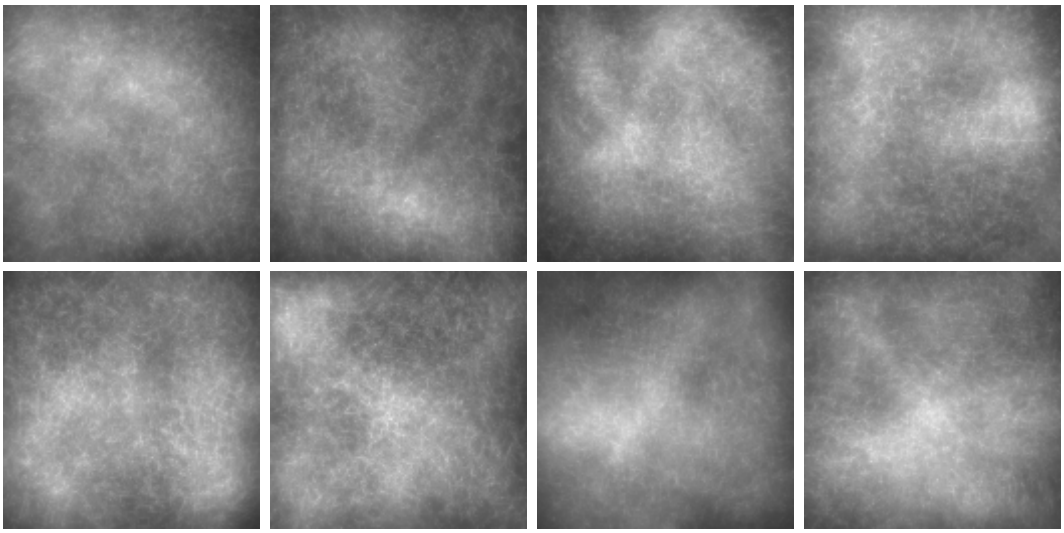

Figure 8.1: Example of clustered lumpy backgrounds (CLBs) used in the observer studies. The images in the bottom row contain the Gaussian signal.

false positives and false negatives, because each test trial in a 2AFC session uses exactly one positive image (containing the synthetically added target) and one negative image (without the target) and the two choices are mutually exclusive. Since the target always has the same shape and location, the user is not required to perform a search and his decision is simplified; for this reason, the time he employs to make each decision is not considered relevant and is neither limited nor accounted for in the evaluation of the test results [4].

The 2AFC technique requires that the user selects one image mandatorily at each trial; if he is not able to detect the target in neither image, he makes a random guess. Since the target is positioned randomly in the left or right image with equal probability, an user who is never able to detect the target and always makes a random guess should obtain approximately a percent correct rate of 50%. The actual detectability of the target can be roughly estimated as follows. We indicate with N the total number of trials (each corresponding to a pair of images) and with n_d the number of trials in which the user is able to detect the target. In these n_d trials, the user will give a correct answer; in the remaining $N - n_\mathrm{d}$ trials, instead, the user will make a random guess and give the correct answer in approximately half the cases. Therefore, the average number of correct decisions, which we shall indicate with n_c, is given by the expression

$$n_\mathrm{c} = n_\mathrm{d} + \frac{N - n_\mathrm{d}}{2} = \frac{N + n_\mathrm{d}}{2}. \tag{8.1}$$

The probability of correct decision is therefore

$$p_\mathrm{c} \triangleq \frac{n_\mathrm{c}}{N} = \frac{1}{2}\left(1 + \frac{n_\mathrm{d}}{N}\right) \triangleq \frac{1 + p_\mathrm{d}}{2}, \tag{8.2}$$

where $p_\mathrm{d} \triangleq n_\mathrm{d}/N$ is the probability of detection. By inverting the expression (8.2), we can estimate the probability of detection p_d as

$$p_\mathrm{d} \approx 2p_\mathrm{c} - 1, \tag{8.3}$$

where the approximation comes from the fact that p_c is estimated by computing the rate of success in a finite number of trials.

The 2AFC method can be easily used to compare two or more image processing algorithms. At each test trial, the test software selects one of the possible algorithms randomly and processes both images using the same parameter settings. The percentage of correct decisions is then computed separately for each algorithm. As previously anticipated, we used this procedure to compare the observer's performance when using a high dynamic range, 16-bit image displayed on the dual layer monitor and a low dynamic range, 8-bit image displayed on a single panel.

The observer studies are currently in progress. Preliminary results have showed that the dual layer presentation mode provides an increase of approximately 8.87% in the probability of correct decision in the detection task we considered. We are planning to conduct a more complete series of studies, with a large number of observers and comparing different characteristics in the presentation modes also regarding the luminance mapping and the parameters in the splitting algorithm.

8.3 Introduction to model observers

The detection of a target in an image can be cast as a *classification* task. Given a set of items, a classification task aims at assigning each item to one category, chosen among two or more predetermined alternatives, depending on some properties of the item itself. In the present case, each item corresponds to an image and the two categories are the "positive" one (with the images containing the target) and the "negative" one (with the images not containing the target). In Section 8.2 we described an experimental procedure, based on the 2AFC method, in which human observers perform a manual classification. The human observers in a classification task can be partially replaced by appropriate software tools, known as *model observers* or *computational observers* [10], which attempt to simulate the human decision by means of statistical estimation techniques. If the model observer is well designed, its use allows to conduct a significantly larger number of test trials compared to human observer studies with a comparable accuracy.

The general structure of a statistical classification method can be described as follows. We consider a training set of N items and measure, for each item, M different properties or *features*. We shall assume for simplicity that each feature is represented by a real number; the training set is therefore represented by a set of N vectors $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N\}$ of M elements, where each vector $\mathbf{x}_i \in \mathbb{R}^M$ represents an item in the training set and each element of the vector represents one of the measured features. We consider in particular a signal detection task which uses two categories: the "positive" one of the items containing the signal and the "negative" one of the items not containing the signal. We shall indicate with $\mathcal{I}_{\text{pos}}$ the set of indices corresponding to the positive items and with $\mathcal{I}_{\text{neg}}$ the set of indices corresponding to the negative items in the training set. If the decision task is well posed and the features are chosen appropriately, it is likely that, when the vectors are plotted as points in $\mathbb{R}^M$, the points corresponding to the positive and negative items tend to gather in two distinct spatial clusters. In this case, it is possible to partition the feature space $\mathbb{R}^M$ in two regions, which we shall indicate with $\mathcal{R}_{\text{pos}}$ and $\mathcal{R}_{\text{neg}}$, containing the points corresponding respectively to the positive and negative items in the training set. After this training phase is completed, a new item can be classified automatically by computing its feature vector and identifying the region $\mathcal{R}_{\text{pos}}$ or $\mathcal{R}_{\text{neg}}$ in which it falls. An exact partitioning is not always possible in general. If the two spatial clusters overlap partially, some positive items $\mathbf{x}_i$ with $i \in \mathcal{I}_{\text{pos}}$ will fall in the region $\mathcal{R}_{\text{neg}}$, and in a similar way, some

negative items $\mathbf{x}_i$ with $i \in \mathcal{I}_{\text{neg}}$ will fall in the region $\mathcal{R}_{\text{pos}}$. The items belonging to the first category are called *false negatives* and the ones belonging to the second category are called *false positives*.

Several techniques have been proposed which estimate the partitions $\mathcal{R}_{\text{pos}}$ and $\mathcal{R}_{\text{neg}}$ of the feature space $\mathbb{R}^M$ using different types of surfaces as a boundary. The most simple techniques perform the partition using a hyperplane and are therefore known as *linear classifiers*. A linear classifier performs the classification of an item $\mathbf{x}_i$ by means of two successive steps. A *decision variable* t_i is computed by taking the scalar product between the feature vector $\mathbf{x}_i$ and a vector $\mathbf{w}$ normal to the hyperplane. We shall assume that the vector $\mathbf{w}$ is oriented towards the region $\mathcal{R}_{\text{pos}}$.

$$t_i \triangleq \mathbf{w}^{\mathrm{T}} \mathbf{x}_i \tag{8.4}$$

The item is then assigned to the positive category if the decision variable takes a value greater than a specific threshold value t_{lim}, or to the negative category otherwise.

The vector $\mathbf{w}$ which describes the orientation of the separating hyperplane is generally computed by means of an appropriate optimization procedure, and different optimality criteria have been proposed. Support Vector Machines [11], for instance, construct the separating hyperplane which maximizes the "margin" between the two data sets; this margin is represented by the distance from the hyperplane to the nearest data point. Barrett et al. [3] proposed to use a linear classifier introduced in 1931 by Harold Hotelling as a model observer. The Hotelling observer computes the hyperplane which maximizes the "signal to noise ratio" (SNR) in the classification, which is defined as follows. We indicate with μ_{pos} the mean value of the decision variable $t \triangleq \mathbf{w}^{\mathrm{T}} \mathbf{x}$ for the positive items in the training set and with σ^2_{pos} its variance. In a similar way, we indicate with μ_{neg} and σ^2_{neg} the mean and variance of the decision variable for the negative items. The SNR is defined as

$$SNR \triangleq \frac{\mu_{\text{pos}} - \mu_{\text{neg}}}{\sqrt{\frac{1}{2}\sigma^2_{\text{pos}} + \frac{1}{2}\sigma^2_{\text{neg}}}}. \tag{8.5}$$

The Hotelling observer computes the hyperplane which maximizes the SNR defined in (8.5). Some computation can show that the means and variances in the expression (8.5) can be computed as

$$\mu_{\text{pos}} = \mathbf{w}^{\mathrm{T}} \bar{\mathbf{x}}_{\text{pos}} \qquad \mu_{\text{neg}} = \mathbf{w}^{\mathrm{T}} \bar{\mathbf{x}}_{\text{neg}} \qquad \sigma^2_{\text{pos}} = \mathbf{w}^{\mathrm{T}} \mathbf{R}_{\text{pos}} \mathbf{w} \qquad \sigma^2_{\text{neg}} = \mathbf{w}^{\mathrm{T}} \mathbf{R}_{\text{neg}} \mathbf{w}, \tag{8.6}$$

where $\bar{\mathbf{x}}_{\text{pos}}$, $\bar{\mathbf{x}}_{\text{neg}}$, $\mathbf{R}_{\text{pos}}$ and $\mathbf{R}_{\text{neg}}$ are the mean values and the covariance matrices of the feature vectors of the positive and negative items in the training set. By substituting the expressions (8.6) into (8.5) and simplifying, we obtain

$$SNR \triangleq \frac{\mathbf{w}^{\mathrm{T}} (\bar{\mathbf{x}}_{\text{pos}} - \bar{\mathbf{x}}_{\text{neg}})}{\sqrt{\frac{1}{2}\mathbf{w}^{\mathrm{T}} (\mathbf{R}_{\text{pos}} + \mathbf{R}_{\text{neg}}) \mathbf{w}}}. \tag{8.7}$$

The vector $\mathbf{w}$ that maximizes the expression (8.7) is

$$\mathbf{w} \propto (\mathbf{R}_{\text{pos}} + \mathbf{R}_{\text{neg}})^{-1} (\bar{\mathbf{x}}_{\text{pos}} - \bar{\mathbf{x}}_{\text{neg}}), \tag{8.8}$$

where the proportionality symbol is justified by the fact that a scaling of the vector $\mathbf{w}$ does not change the orientation of the normal plane.

In the specific task of the detection of a target in a radiographic image, the training set for the model observer consists of a collection of test images – either natural or synthetic – for which a ground truth classification is available. This classification is typically manual, if natural images are used, or automatic, if the test target is added synthetically. Ideally, the feature vector for each image could be made of the pixel values directly. However, this approach is only applicable in practice to images with a very small size; for large images, the expression (8.8) would involve the solution of a very large linear system and have a prohibitive computational cost. It is therefore preferable to use a reduced set of features, and the choice of the features which allow a significant dimensionality reduction while preserving the accuracy of the decision is still an active research topic. For instance, channelized Hotelling observers (CHOs) consider appropriate weighted averages of the image pixels.

In some cases, the model observers can surpass the performance of human observers. Techniques have been proposed [55] which incorporate models of the human visual system (HVS) in the selection of the channels, in order to simulate as accurately as possible the performance of a human observer in the quality assessment of medical imaging systems. We are currently attempting to design an observer which incorporates a model of the behavior of the HVS in presence of high dynamic range images, in order to build an useful tool in the performance evaluation of a dual layer LCD display.

Bibliography

[1] M. Ashikhmin. A tone mapping algorithm for high contrast images. In *13th Eurographics Workshop on Rendering*, pages 1–11, June 2002.

[2] Tunç Ozan Aydn, Rafał Mantiuk, Karol Myszkowski, and Hans-Peter Seidel. Dynamic range independent image quality assessment. In *Proc. SIGGRAPH 2008*.

[3] H. H. Barrett, J. Yao, J. P. Rolland, and K. J. Myers. Model observers for assessment of image quality. In *Proceedings of the National Academy of Science, USA*, volume 90, pages 9758–9765, 1993.

[4] Harrison H. Barrett and Kyle J. Myers. *Foundations of Image Science*. John Wiley & Sons, October 2003.

[5] Peter G. Barten. Physical model for the contrast sensitivity of the human eye. In *Proceedings of SPIE*, volume 1666, pages 57–72, August 1992.

[6] Peter G. Barten. Spatio-temporal model for the contrast sensitivity of the human eye and its temporal aspects. In *Proceedings of SPIE*, volume 1913, pages 2–14, September 1993.

[7] Peter G. J. Barten. *Contrast Sensitivity of the Human Eye and its Effects on Image Quality*. SPIE Optical Engineering Press, 1999.

[8] Bryce E. Bayer. Color imaging array, July 1976. U.S. Patent No. 3,971,065.

[9] Dimitri P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 2nd edition edition, September 1999.

[10] Jacob Beutel, Harold L. Kundel, and Richard L. Van Metter. *Handbook of Medical Imaging – Volume 1. Physics and Psychophysics*. SPIE Press, February 2000.

[11] B. E. Boser, I. M. Guyon, and V. N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 144–152. ACM Press, 1992.

[12] Achi Brandt. Multi-level adaptive solutions to boundary-value problems. *Mathematics of Computation*, 31(138):333–390, April 1977.

[13] Tony F. Chan, Stanley Osher, and Jianhong Shen. The digital TV filter and nonlinear denoising. *IEEE Transactions on Image Processing*, 10:231–241, 2001.

[14] CIE Technical Report No. 19/2. *An analytic model for describing the influence of lighting parameters upon visual performance – Vol. 1: Technical foundations*, 1981.

[15] Scott Daly. *The Visible Differences Predictor: An algorithm for the assessment of image fidelity*, pages 179–206. MIT Press, Cambridge, MA, 1993.

[16] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH 97*, August 1997.

[17] F. Drago, K. Myszkowski, T. Annen, and N. Chiba. Adaptive logarithmic mapping for displaying high contrast scenes. In *EUROGRAPHICS 2003*, volume 22, 2003.

[18] Frédo Durand and Julie Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. In *SIGGRAPH 2002*, 2002.

[19] Raanan Fattal, Dani Lischinski, and Michael Werman. Gradient domain high dynamic range compression. In *ACM SIGGRAPH 2002*, July 2002.

[20] J. A. Ferwerda, S. N. Pattaniak, P. Shirley, and D. P. Greenberg. A model of visual adaptation for realistic image synthesis. In *Computer Graphics*, pages 249–258, 1996.

[21] Jonathan A. Frankle and John J. McCann. Method and apparatus for lightness imaging, May 1983. U.S. Patent No. 4,384,336.

[22] Brian Funt, Florian Ciurea, and John McCann. Retinex in Matlab. *Journal of Electronic Imaging*, 13(1):48–57, January 2004.

[23] Gabriele Guarnieri, Luigi Albani, and Giovanni Ramponi. Image-splitting techniques for a dual-layer high dynamic range LCD display. *Journal of Electronic Imaging*, 17(4):043009–1–9, Oct-Dec 2008.

[24] Gabriele Guarnieri, Luigi Albani, and Giovanni Ramponi. Image splitting techniques for a dual layer high dynamic range LCD display. In *Proc. Medical Imaging 2008 – Image Perception, Observer Performance, and Technology Assessment*, volume 6917, page 69170M. SPIE, March 2008.

[25] Gabriele Guarnieri, Luigi Albani, and Giovanni Ramponi. Minimum-error splitting algorithm for a dual layer LCD display - Part I: Background and theory. *Journal of Display Technology*, 4(4):383–390, Dec 2008.

[26] Gabriele Guarnieri, Luigi Albani, and Giovanni Ramponi. Minimum-error splitting algorithm for a dual layer LCD display - Part II: Implementation and results. *Journal of Display Technology*, 4(4):391–397, Dec 2008.

[27] Gabriele Guarnieri, Sergio Carrato, and Giovanni Ramponi. Nonlinear mapping for dynamic range compression in digital images. In *Proc. Electronic Imaging 2007 – Image Processing: Algorithms and Systems V*, volume 6497, page 64970M. SPIE, 2007.

[28] Gabriele Guarnieri, Sergio Carrato, and Giovanni Ramponi. Psychophysical evaluation and tuning of a simple tone mapping operator. In *Proc. NSIP 2007 – International Workshop on Nonlinear Signal and Image Processing*, pages 122–128, September 2007.

[29] Gabriele Guarnieri, Stefano Marsi, and Giovanni Ramponi. Fast bilateral filter for edge-preserving smoothing. *Electronics Letters*, 42(7), March 2006.

[30] Gabriele Guarnieri, Giovanni Ramponi, Silvio Bonfiglio, and Luigi Albani. Nonlinear mapping of the luminance in dual-layer high dynamic range displays. In *Proc. Electronic Imaging 2009 – Image Processing: Algorithms and Systems VII*, volume 7245, page 72450D. SPIE, 2009.

[31] J. Guild. The colorimetric properties of the spectrum. *Philosophical Transactions of the Royal Society of London*, A230:149–187, 1931.

[32] Wolfgang Hackbusch. *Multigrid methods and applications*. Springer, Berlin, 1985.

[33] Francis B. Hildebrand. *Methods of applied mathematics*. Dover, New York, 2nd edition edition, 1992.

[34] Berthold K. P. Horn. Determining lightness from an image. *Computer Graphics and Image Processing*, 3:277–299, 1974.

[35] Gaetano Impoco, Anna Ukovich, and Giovanni Ramponi. Exploiting the local contrast to evaluate image dynamics reduction algorithms. In *Proc. 8th COST 276 Workshop*, May 2005.

[36] International Color Consortium. *Specification ICC.1:2004-10 (Profile version 4.2.0.0). Image technology colour management – Architecture, profile format, and data structure*, 2004.

[37] International Telecommunication Union. *Recommendation ITU-R BT.709-4: Parameter Values for the HDTV Standards for Production and International Programme Exchange*, 2000.

[38] H. J. Dartnall J. K. Bowmaker. Visual pigments of rods and cones in a human retina. *The Journal of Physiology*, 298(1):501–511, 1980.

[39] James Keener and James Sneyd. *Mathematical Physiology*. Springer, 1998.

[40] Ron Kimmel, Michael Elad, Doron Shaked, Renato Keshet, and Irwin Sobel. A variational framework for Retinex. *International Journal of Computer Vision*, 52(1):7–23, 2003.

[41] Ron Kimmel, Doron Shaked, Michael Elad, and Irwin Sobel. Space-dependent color gamut mapping: A variational approach. *IEEE Transactions on Image Processing*, 14(6), June 2005.

[42] K. M. Lam. Metamerism and colour constancy, 1985. PhD thesis.

[43] E. H. Land. The Retinex theory of color vision. *Scientific American*, 237(6):108–128, December 1977.

[44] E. H. Land and J. J. McCann. Lightness and Retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, January 1971.

[45] R. D. Luce and W. Edwards. The derivation of subjective scales from just noticeable differences. *Psychological Review*, 65(4):222–237, July 1958.

[46] D. L. MacAdam. Visual sensitivities to color differences in daylight. *Journal of the Optical Society of America*, 32(5):247–274, May 1942.

[47] Rafał Mantiuk, Karol Myszkowski, and Hans-Peter Seidel. Visible difference predicator for high dynamic range images. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, pages 2763–2769, 2004.

[48] Rafał Mantiuk, Karol Myszkowski, and Hans-Peter Seidel. A perceptual framework for contrast processing of high dynamic range images. In *APGV 05: Proc. 2nd symposium on Appied perception in graphics and visualization*, pages 87–94, 2005.

[49] S. Marsi, G. Ramponi, and S. Carrato. Image contrast enhancement using a recursive rational filter. In *IEEE IST 2004 International Workshop on Imaging Systems and Techniques*, May 2004.

[50] John McCann. Lessons learned from Mondrians applied to real images and color gamuts. In *Proc. IS&T/SID Seventh Color Imaging Conference*, pages 1–8, 1999.

[51] P. Moon and D. Spencer. The visual effect of nonuniform surrounds. *Journal of the Optical Society of America*, 35(3):233–248, 1945.

[52] National Electrical Manufacturers Association. *Digital Imaging and Communications in Medicine (DICOM). Part 14: Grayscale Standard Display Function*, 2007.

[53] Alan V. Oppenheim, Ronald W. Schafer, and Jr. Thomas G. Stockham. Nonlinear filtering of multiplied and convolved signals. In *Proceedings of the IEEE*, volume 56, August 1968.

[54] Sylvain Paris and Frédo Durand. A fast approximation of the bilateral filter using a signal processing approach. In *Proc. European Conference on Computer Vision*, pages 568–580, 2006.

[55] Subok Park, Aldo Badano, Brandon D. Gallas, and Kyle J. Myers. A contrast-sensitive channelized-Hotelling observer to predict human performance in a detection task using lumpy backgrounds and Gaussian signals. In *Proc. Medical Imaging 2007 – Image Perception, Observer Performance, and Technology Assessment*, volume 6515, page 65150V. SPIE, 2007.

[56] Carinna Parraman, Alessandro Rizzi, and John J. McCann. Color appearance and color rendering of HDR scenes: an experiment. In *Color Imaging XIV: Displaying, Processing, Hardcopy, and Applications*, volume 7241, page 72410R, 2009.

[57] D. W. Peaceman and H. H. Rachford, Jr. The numerical solution of parabolic and elliptic differential equations. *Journal of the Society for Industrial and Applied Mathematics*, 3(1):28–41, March 1955.

[58] Perry A. Penz. Stacked electro-optic display, December 1982. U.S. Patent No. 4,364,039.

[59] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7), July 1990.

[60] Charles Poynton. *A Technical Introduction to Digital Video*. John Wiley & Sons, 1996.

[61] G. Ramponi. The rational filter for image smoothing. *IEEE Signal Processing Letters*, 3(3):63–65, March 1996.

[62] L. Roberts. *Machine Perception of 3-D Solids*. MIT Press, 1965.

[63] Mark A. Robertson, Sean Borman, and Robert L. Stevenson. Estimation-theoretic approach to dynamic range enhancement using multiple exposures. *Journal of Electronic Imaging*, 12(2):219–228, April 2003.

[64] Yousef Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, 2003.

[65] David Salomon. *Curves and Surfaces for Computer Graphics*. Springer-Verlag, 2006.

[66] C. Schlick. Quantization techniques for visualization of high dynamic range pictures. In P. Shirley, G. Sakas, and S. Müller, editors, *Photorealistic Rendering Techniques*, pages 7–20. Springer-Verlag, 1994.

[67] H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. A. Whitehead, M. Trentacoste, A. Gosh, and A. Vorozcovs. High dynamic range display systems. In *Proc. ACM SIGGRAPH 2004*, August 2004.

[68] Doron Shaked and Renato Keshet. Robust recursive envelope operators for fast Retinex. *Technical Note HPL-2002-74 (R.1), HP Laboratories Israel*, March 2004.

[69] Alvy Ray Smith. Color gamut transform pairs. In *SIGGRAPH 78 Conference Proceedings*, pages 12–19, August 1978.

[70] Robert Sobol. Improving the Retinex algorithm for rendering wide dynamic range photographs. *Journal of Electronic Imaging*, 13(1):65–74, January 2004.

[71] J. T. G. Stockham. Image processing in the context of a visual model. In *Proceedings of the IEEE*, volume 60, pages 828–842, 1972.

[72] Michael Stokes, Matthew Anderson, Srinivasan Chandrasekar, and Ricardo Motta. *A Standard Default Color Space for the Internet: sRGB, Version 1.10*, November 1996.

[73] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proc. Sixth International Conference on Computer Vision*, January 1998.

[74] M. Trentacoste, W. Heidrich, L. Whitehead, H. Seetzen, and G. Ward. Photometric image processing for high dynamic range displays. *Journal of Visual Communication and Image Representation*, 18(5):439–451, October 2007.

[75] Jack Tumblin and Greg Turk. LCIS: a boundary hierarchy for detail-preserving contrast reduction. In *SIGGRAPH 99*, pages 83–90, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.

[76] H. M. Visser, J. J. W. M. Rosink, N. Raman, and R. Rajae-Joordens. Invited paper: Tuning LCD displays to medical applications. In *EuroDisplay 2005*, pages 74–77, 2005.

[77] J. Vos. Disability glare - a state of the art report. *CIE Journal*, 3(2):39–53, 1984.

[78] G. Ward Larson, H. Rushmeier, and C. Piatko. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, 1997.

[79] W. D. Wright. A re-determination of the trichromatic coefficients of the spectral colours. *Transactions of the Optical Society*, 30:141–164, 1928.

[80] Günter Wyszecki and W. S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulae (2nd edition)*. John Wiley & Sons, 1982.

[81] Z. Xie and T. G. Stockham, Jr. Toward the unification of three visual laws and two visual models in brightness perception. *IEEE Transactions on Systems, Man and Cybernetics*, 19(2):379–387, Mar/Apr 1989.

Acknowledgments

The research presented in this Thesis was made possible by the financial support of different Institutions. I gratefully acknowledge a PhD studentship grant of the University of Trieste; this work was also partially supported by a grant of the Regione Friuli Venezia Giulia within the "Eladin" and "Eladin II" projects. The research on the dual layer LCD displays for medical applications was conducted in cooperation with the FIMI-Philips company in Saronno. Besides building the prototypes, granting access to their research laboratories and sharing their technical knowledge, FIMI-Philips provided a significant financial contribution to this study. I particularly appreciate their decision of publishing the research results, thus making possible a significant portion of this Thesis and of my publications cited in the Bibliography.

The evaluation and tuning of image processing algorithms by means of psychophysical experiments was partially inspired by a visit to the Image Processing Research Laboratory of the University of California, Santa Barbara. In particular I wish to thank Professor Sanjit K. Mitra, at the Department of Electrical and Computer Engineering, and Professor John M. Foley, at the Department of Psychology, for kindly giving me some valuable advice on this research subject. I wish to thank my friends and colleagues who dedicated some of their time taking part in the experiments. I am also grateful to Prof. Walter Gerbino and Dr. Carlo Fantoni, at the Department of Psychology of the University of Trieste, for their cooperation in organizing psychophysical experiments in our laboratory, and to my friend and former colleague Livio Tenze for his help in the design of the interface for the program used in the experiments.

The research on a dual layer LCD display for medical applications was partially conducted during a visit at the Center for Devices and Radiological Health of the U.S. Food and Drug Administration. I wish to thank Dr. Aldo Badano, Leader of the Imaging Physics Laboratory in the Division of Imaging and Applied Mathematics, for his kind hospitality and for the helpful suggestions and discussions in the design of the splitting algorithms and in the test programs for observer studies. I am grateful to Dr. Anindita Saha, for her assistance in the use of the IL1700 radiometer, and to Dr. Subok Park, for her valuable cooperation in the design of computational observers.

I wish to thank my supervisors for the helpful discussions and encouragement throughout this work. I am grateful in particular to Prof. Giovanni Ramponi, Sergio Carrato and Stefano Marsi at the University of Trieste, Luigi Albani and Silvio Bonfiglio at FIMI-Philips and Dr. Aldo Badano at the US Food and Drug Administration. I thank the reviewers of the journals and conferences in which I published my results for their helpful and constructive comments. Finally, I am grateful to all my colleagues and coworkers for the valuable professional and human cooperation we built during these years.