



National Library
of Canada

Acquisitions and
Bibliographic Services Branch

395 Wellington Street
Ottawa, Ontario
K1A 0N4

Bibliothèque nationale
du Canada

Direction des acquisitions et
des services bibliographiques

395, rue Wellington
Ottawa (Ontario)
K1A 0N4

Your file *Votre référence*

Our file *Notre référence*

NOTICE

The quality of this microform is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

If pages are missing, contact the university which granted the degree.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us an inferior photocopy.

Reproduction in full or in part of this microform is governed by the Canadian Copyright Act, R.S.C. 1970, c. C-30, and subsequent amendments.

AVIS

La qualité de cette microforme dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de qualité inférieure.

La reproduction, même partielle, de cette microforme est soumise à la Loi canadienne sur le droit d'auteur, SRC 1970, c. C-30, et ses amendements subséquents.

UNIVERSITY OF ALBERTA

The Local K-means Algorithm for Colour Image Quantization.

BY

Oleg Verevka



A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Master of Science.

DEPARTMENT OF COMPUTING SCIENCE

Edmonton, Alberta
Fall 1995



National Library
of Canada

Acquisitions and
Bibliographic Services Branch

395 Wellington Street
Ottawa, Ontario
K1A 0N4

Bibliothèque nationale
du Canada

Direction des acquisitions et
des services bibliographiques

395, rue Wellington
Ottawa (Ontario)
K1A 0N4

Your file *Voire référence*

Our file *Notre référence*

THE AUTHOR HAS GRANTED AN IRREVOCABLE NON-EXCLUSIVE LICENCE ALLOWING THE NATIONAL LIBRARY OF CANADA TO REPRODUCE, LOAN, DISTRIBUTE OR SELL COPIES OF HIS/HER THESIS BY ANY MEANS AND IN ANY FORM OR FORMAT, MAKING THIS THESIS AVAILABLE TO INTERESTED PERSONS.

L'AUTEUR A ACCORDE UNE LICENCE IRREVOCABLE ET NON EXCLUSIVE PERMETTANT A LA BIBLIOTHEQUE NATIONALE DU CANADA DE REPRODUIRE, PRETER, DISTRIBUER OU VENDRE DES COPIES DE SA THESE DE QUELQUE MANIERE ET SOUS QUELQUE FORME QUE CE SOIT POUR METTRE DES EXEMPLAIRES DE CETTE THESE A LA DISPOSITION DES PERSONNE INTERESSEES.

THE AUTHOR RETAINS OWNERSHIP OF THE COPYRIGHT IN HIS/HER THESIS. NEITHER THE THESIS NOR SUBSTANTIAL EXTRACTS FROM IT MAY BE PRINTED OR OTHERWISE REPRODUCED WITHOUT HIS/HER PERMISSION.

L'AUTEUR CONSERVE LA PROPRIETE DU DROIT D'AUTEUR QUI PROTEGE SA THESE. NI LA THESE NI DES EXTRAITS SUBSTANTIELS DE CELLE-CI NE DOIVENT ETRE IMPRIMES OU AUTREMENT REPRODUITS SANS SON AUTORISATION.

ISBN 0-612-06570-7

Canada

UNIVERSITY OF ALBERTA

RELEASE FORM

NAME OF AUTHOR: Oleg Verevka

TITLE OF THESIS: The Local K-means Algorithm for Colour Image Quantization.

DEGREE: Master of Science

YEAR THIS DEGREE GRANTED: 1995

Permission is hereby granted to the University of Alberta Library to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only.

The author reserves all other publication and other rights in association with the copyright in the thesis, and except as hereinbefore provided neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatever without the author's prior written permission.


(Signed) *Oleg Verevka*
Oleg Verevka
45^a , Ivan Franko str, ap.18
Odessa , 270049
Ukraine

Date: *Aug 15, 95*

UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES AND RESEARCH

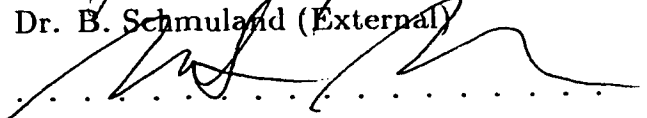
The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled **The Local K-means Algorithm for Colour Image Quantization**. submitted by Oleg Verevka in partial fulfillment of the requirements for the degree of Master of Science.



Dr. J. Buchanan (Supervisor)



Dr. B. Schmuland (External)



Dr. M. Green (Examiner)

Dr. J. Culberson (Chair)

Date: *Aug 15, 95*

Abstract

One of the most important goals of computer graphics is the reliable display of synthesized images. This thesis addresses the problem of image display on computer graphics hardware with a limited set of simultaneous colours.

The main emphasis of this work is colour image quantization. Quantization is the process of representing an image with a small number of well selected colours. The thesis gives a detailed survey of previous image quantization techniques. The existing algorithms are divided into two classes — pre- and post-clustering. Merits of both approaches are discussed. In addition this document presents optimization techniques for the nearest neighbour search. These methods may be used to speed up the iteration process of the post-clustering algorithms.

A new colour quantization approach is presented — the local K-means algorithm. It is an iterative post-clustering technique that approximates an optimal palette using multiple subsets of image pixels. The local K-means procedure is compared with previous quantization methods. The new algorithm is able to generate a high quality palette significantly faster than other quantization techniques. The application of the local K-means algorithm to quantization of multiple images in windows systems is addressed. The algorithm takes into account previously allocated colours in the shared colour map, thus the quantization accuracy may be improved.

Acknowledgements

This work is a journey into unknown, and when this journey is completed I would like to thank all those without whom this work would not be possible:

My family — my dear parents, brother — for giving me the strength and will to go. Victoria for staying together with me in my victories and failures.

John Buchanan, my research advisor, for suggesting the direction and keeping the door opened when the path was not clear, for encouraging me to stay on the road and to see the new goals.

My examination committee, for taking the time to read this work, and for valuable comments and suggestions.

The graphics group at the University of Alberta, the administration and support staff for funding my studies and providing computing facilities.

To all my friends, who made this journey enjoyable.

Contents

1	Introduction.	1
1.1	Frame buffer displays	2
1.2	The scope of this thesis	2
1.2.1	<i>Pre-</i> and <i>post-</i> processing	3
1.2.2	Quantization	4
1.3	The “road map”	5
2	Formulation of the Problem.	6
2.1	Quantization is an optimization problem.	6
2.2	Quantization errors	8
2.2.1	Measures of Colour Approximation	8
2.2.2	Context-dependent colour mapping	9
2.3	Approaches to colour quantization	11
2.4	Algorithm evaluation methodology	14
2.4.1	Visualization of quantization errors	15
3	Pre-clustering	17
3.1	Median cut	17
3.2	Variance-based algorithm.	20
3.3	Center cut	22
3.4	Octree	22
3.5	Principal axis splitting quantization.	23
3.6	Wrapup	27
4	Post-clustering algorithms	29
4.1	K-means algorithm	29
4.2	Kohonen self-organizing maps.	30
4.3	Simulated annealing	33
4.4	Motivation for the future research	34
5	Fast nearest neighbour search.	35
5.1	L_α -norm	36
5.2	Optimization techniques	36
5.2.1	Calculation of the partial sum	37
5.2.2	Sorting on one coordinate	37

5.2.3	Nearest neighbour distance (NND)	38
5.2.4	Experiments in colour mapping	38
6	Local K-means for colour quantization.	41
6.1	The local K-means algorithm	41
6.1.1	The adaptation procedure.	42
6.1.2	The adaptation parameter	42
6.1.3	The choice of the initial palette.	43
6.1.4	Input data.	44
6.2	Quantization results.	45
6.2.1	Loss of colour information.	45
6.2.2	Artificial contouring.	48
6.2.3	Compromises	50
6.3	Performance optimization.	54
6.4	Wrapup	56
7	Multiple Image Quantization with Local K-means Algorithm.	57
7.1	Colour quantization of multiple images	57
7.2	Modified local K-means quantization	58
7.3	Experiments with quantization for windows systems	59
7.4	Discussion	60
	Bibliography	62

List of Figures

2.1	The nature of artificial contouring in quantization	10
2.2	Error propagation in the Floyd-Steinberg algorithm	11
2.3	Application of dithering to colour mapping	12
2.4	Peano scan encoding of the colour space	13
2.5	The test images.	16
3.1	Partitioning of the space: <i>Median cut</i>	18
3.2	Median cut quantization test.	19
3.3	Partitioning of the space: <i>Variance based quantization</i>	20
3.4	Variance-based quantization test	21
3.5	Partitioning of the space: <i>Octree quantization</i>	23
3.6	Octree quantization test	24
3.7	Inertial cut quantization test	26
3.8	Partitioning of the space: <i>Principal multilevel quantization</i>	27
3.9	Voronoi regions of the median cut partitioning	28
4.1	K-means adaptation scheme	30
4.2	SOM adaptation scheme	31
4.3	SOM quantization test	32
4.4	Simulated annealing quantization	33
5.1	Partial sum criteria	38
5.2	Sorting by one coordinate	39
5.3	Nearest neighbour distance	39
6.1	Quantization error images “Kiss”	46
6.2	“Kiss”: quantization to 16 colours	47
6.3	Quantization error images “Pool balls”	49
6.4	The original portrait of Anita Yuen	51
6.5	Quantization to 32 colours (the portrait of Anita Yuen)	52
6.6	Quantization to 32 colours with Floyd - Steinberg error diffusion	53
6.7	Complexity analysis	55

List of Tables

5.1	Application of the $L_{\alpha=1/2}$ to the colour mapping.	37
5.2	Performance of the colour mapping optimization techniques.	40
6.1	Quantization errors: image “Kiss”, 16 colours.	45
6.2	Quantization errors: image “Pool Balls”, 256 colours.	48
6.3	Quantization errors: the Anita Yuen portrait, 32 colours	51
6.4	Execution time in seconds for colour map selection	54
7.1	Quantization errors for local K-means algorithm applied in windowing systems, image “Kiss”	59

Chapter 1

Introduction.

Visual information is extremely important in our perception of the world. This is why computing applications are becoming more and more *visual*. Modern software incorporates various forms of imagery. Image input devices are able to accurately digitize real world scenes, photographs, works of art, etc. Computer graphics researchers have developed highly photo-realistic image synthesis techniques. Such rendering algorithms are based on methods for the faithful representation of geometric objects, and mathematical models of the physical behavior of light. Unfortunately, the effort in image acquisition and synthesis may be lost due to the limitations of the display hardware.

Alain Fournier said “*Computer Graphics is computer science you can see* ” [Fou94]. The objective of this work is to make *visible* the already generated image information. In particular, the goal of this research is to display faithfully full colour images using relatively scarce resources of output devices.

The complete solution of the image display problem is far beyond the scope of this thesis. This work is targeted to the reproduction of images on hardware with a limited set of simultaneous colours, mainly frame buffer displays.

1.1 Frame buffer displays

The images are often acquired or synthesized using a 24 bit representation of colour values. Such a representation allows more than 16 million different colours. This is sufficient for most applications since a typical image of a real scene contains only a few hundred thousand distinct colours. In order to be displayed the image needs to be stored in the frame buffer. Unfortunately the 24 bit representation requires a substantial amount of video memory. Commonly used low cost hardware solves the problem by the reduction of the simultaneous colours. The colours displayable at one time constitute the “palette”. The frame buffer stores not the original 24-bit values of image pixels but the index of the corresponding colour in the palette. Thus each pixel is represented by 4 or 8 bits — an index into 16 or 256 colours in the palettes.

Clearly, for the best image display the palette needs to be carefully chosen. The mapping of the original image to the selected palette must ensure the least possible distortion. These issues constitute the problem of colour reduction. Thus, the palette selection and colour mapping problems are the main topic of this thesis.

Even though 24-bit graphics hardware is becoming more common, colour reduction maintains its practical value. It lessens space requirements for storage of image data and reduces transmission bandwidth requirements in multimedia applications.

1.2 The scope of this thesis

The objective of research in colour reduction is to minimize the perceived distortion in the resulting image. It is often the case that the quality of the colour reduction procedure is judged by the appearance of the image on hardware devices with different colour capabilities. In order to be successful the reduction procedure must account for the specifics of the target output hardware. Therefore the issues of colour reproduction should be considered.

The colour reproduction problem has been studied in the cases when an image is to be displayed by hardware of different media: printer, computer monitor, etc. The

goal is to minimize the difference in various representations of the same image.

The author suggests that a full scale colour reduction should be “wrapped” by the *pre-* and *post-*processing operations that allow faithful colour reproduction:

1. Preprocessing.

- (a) Representation of the image in device independent form.
- (b) Device-to-device gamut mapping.
- (c) De-saturation and Over-saturation corrections.
- (d) Gamma correction.
- (e) Representation of the image in device dependent form.

2. Colour quantization.

- (a) Selection of colour palette
- (b) Mapping of the original image into the new palette

3. Conversion of the image into the colour space of the output device

1.2.1 *Pre- and post-processing*

The pre-processing step is needed to accommodate the differences in physical properties of the image acquisition process and the output device. The pre-processing is also needed when the objective of the colour reduction is to minimize the difference between the image displayed on the full colour monitor and the 4 or 8 bit frame buffer.

Even though the monitors share similar technology they may use different phosphors to represent primary colours. Moreover, the optical properties of the device are not constant in time. The phosphors fade with use. These differences lead to different sets of reproducible colours. Therefore it is very likely that the same image will appear to be distorted when displayed on two different monitors.

The details of the pre-processing steps can be found in colour reproduction literature ([SCB88], [Hal89], [Yul67]). In this thesis the experiments were carried out on 24 bit colour hardware. The original and the modified images are displayed on the

same monitor and reproduced on the same printer. Therefore the gamut mapping, tone, and saturation corrections were not required.

The only pre-processing that may be needed is the image conversion into device independent colour space. The advantages of the CIE uniform spaces will be explored further in the thesis. Thus the post-processing is a conversion of the image into monitor's RGB space. The methods of this conversion are described in [FvDFH90] (pp. 584-594).

1.2.2 Quantization

Colour quantization is an important step in the colour reduction process. One of the possible definitions is found in [Hec82]:

Quantization is the process of assigning representation values to ranges of input values ...

Colour image quantization is the process of selecting a set of colours to represent the colour gamut of an image, and computing the mapping from colour space to representative colours.

Colour quantization is a *lossy image compression* operation. Other approaches to image compression are developed by the signal processing community. The main emphasis of their research is to achieve high compression rates while maintaining minimum distortion. The image palettes usually do not change. The colour quantization is targeted to the specifics of the frame buffers. Thus the build-in hardware lookup tables can be used for fast display.

The approaches to colour reduction problem differ in the palette selection strategy. The palette entries can be fixed for all the displayed images. This technique is often used in applications when the entire image is not available — preview to a ray-tracer, World Wide Web browsers, etc. Dithering methods are often applied in this case.

The other approach is to select a different palette for each displayed image. Such an algorithm takes into account the statistical distribution of colours in the input image. The thesis is focussed on this class of quantization methods.

1.3 The “road map”

The current work is a study of the colour reduction — the colour quantization problem. The mathematical formulation of this problem is given in Chapter 2. The goal of quantization research is to find the best possible palette and to map the colours of the original image to the selected colours.

This document presents a survey of palette selection schemes developed in the last twenty years. The author describes the new local K-means algorithm. This method favorably compares with previous techniques in accuracy, speed and resource requirements.

This research also deals with fast colour mapping. Several methods of optimization are proposed. Special attention is paid to the minimization of quantization artifacts.

Chapter 2

Formulation of the Problem.

In this chapter we formulate the colour quantization problem. The objective of colour reduction is defined mathematically in terms of the quantization error. Colour mapping techniques are presented along with a brief overview of the early approaches to palette selection.

2.1 Quantization is an optimization problem.

Let c_i be a 3-dimensional vector in one of the colour spaces (CIE Lu^*v^* , HSV, RGB, etc.). The set $C = \{c_i, i = 1, 2 \dots N\}$ is the set of all colours in the full colour image I . A quantized image \bar{I} is represented by a set of K colours $\bar{C} = \{\bar{c}_j, j = 1, 2 \dots K\}, K \ll N$. The quantization process is therefore a mapping: $q : C \rightarrow \bar{C}$ that substitutes each original colour by a colour from the palette.

The goal of quantization is to make the perceived difference between the original image and its quantized representation as small as possible. Hence the colour mapping should substitute every colour of the original image I by the closest colour from the set \bar{C} . Therefore the quantization operator is commonly expressed as follows:

$$q(c) = \bar{c}_k : d(c, \bar{c}_k) = \min_{\bar{c}_j \in \bar{C}} d(c, \bar{c}_j) \quad (2.1)$$

where $d(c_i, c_j)$ is a perceptually meaningful colour distance between colours c_i and c_j . The definition of such a metric is a challenge in itself. The current results show that in

the case of the CIE Lu^*v^* space, the Euclidean distance can reasonably approximate the perceived difference of colours. Thus for the future discussion we will assume that the Euclidean norm is used for the colour mapping operator.

The objective of colour quantization research is to find the best possible palette. Let $\mathcal{E}_{\bar{C}}$ be a measure of image distortion for a palette \bar{C} . We define an optimal palette to be a set of colours \bar{C}^* that best approximates the original image. In other words the optimal palette minimizes the given norm \mathcal{E} :

$$\mathcal{E}_{\bar{C}^*} < \mathcal{E}_{\bar{C}} \quad (2.2)$$

Thus the colour quantization can be viewed as an approximation problem or a vector quantization problem.

On the other hand, the optimal quantization can be formulated as an optimal space partitioning. The quantization process defines a set of clusters in the colour space. We say that colours of the original image mapped into the same palette entry in the output image belong to the same cluster:

$$c \in S_k : \quad q(c) = \bar{c}_k \quad (2.3)$$

The collection of these clusters represent the Voronoi tessellation of the colour space. The partitioning is said to be *optimal* if the following measure is minimized [Wu92a]:

$$\mathcal{E}(S_1, S_2 \dots S_K) = \sum_{k=1}^K \sum_{c \in S_k} \|c - \bar{c}_k\| \quad (2.4)$$

The K clustering problem is known to be NP-complete for variable K ([Bru77], [GJW82], [Wu92a]). Consequently, any practical solution of such a large scale optimization problem will necessarily be heuristic and approximate. It is worth noting that in practice the globally optimal solution is usually not necessary ([WPW90], [FO89]). The nearly optimal partitioning is often sufficient for a good image display.

2.2 Quantization errors

Human vision is an extremely complicated and not yet fully understood process. It is very difficult to formulate a definite solution to image quantization in terms of *perceived* image quality. In fact, there is no good objective criterion available for measuring the perceived image similarity.

At least two typical artifacts are often visible as a result of quantization:

- the colour shift, the loss of colour variety and contrast;
- artificial contouring in the smooth areas of the image.

In this thesis the merits of quantization algorithms are evaluated on the basis of these two artifacts. This section presents mathematical measures of quantization accuracy.

2.2.1 Measures of Colour Approximation

The fidelity of colour approximation is often the only measure of quantization distortion. In fact the *optimal partitioning* is defined only as the most accurate representation of colours.

In the quantization literature it is common to use image dependent distortion measures [Hec82, WPW90, Wu92a]. Let an image I be an array of M pixels (x, y) , then $c_{(x,y)}$ is the colour of each image pixel. The average quantization distortion per pixel can be defined as follows:

$$\mathcal{E} = \frac{1}{M} \sum_{(x,y) \in I} \|c_{(x,y)} - q(c_{(x,y)})\|. \quad (2.5)$$

Even though the average distortion measure \mathcal{E} can give a reasonable estimate of a perceived image difference, it can also be very misleading (see [WPW90]). Colours of the original image are often non-uniformly distributed in the colour space. Thus significant image information is carried by some distinct but “rare” colours (e.g. specular highlights). If a quantization algorithm approximates the more popular colours, the average distortion might be small, but the “rare” colours of the original will be lost.

In this thesis the approximation accuracy is evaluated by a combination of distortion measures. In addition to the quantization error \mathcal{E} we define the average colour distortion:

$$\epsilon = \frac{1}{N} \sum_{i=1}^N \|c_i - q(c_i)\|, \quad (2.6)$$

and the standard deviation of distortion per pixel:

$$\sigma = \sqrt{\frac{\sum_{(x,y) \in I} (\|c_{(x,y)} - q(c_{(x,y)})\| - \mathcal{E})^2}{M - 1}}. \quad (2.7)$$

Small values of ϵ guarantee that a quantization process accurately represents colours of the original image. However the human visual system is not able to determine the absolute value of a colour. It is more sensitive to colour variations. A quantization algorithm that produces small values of σ introduces almost equal colour distortion to every pixel. Therefore the minimization of the standard deviation of distortion, σ , helps us to preserve variations of colours in the quantized image.

It should be noted that these error measures have a significant limitation. The colour context is important in our colour interpretation. Despite the fact that \mathcal{E} , ϵ , and σ are image dependent measures, they treat each pixel independently. Spatial correlation among colours is not taken into account. Balasubramanian and Allebach in [BA91] attempted to account for the colour context by a pre-quantization step. Unfortunately the technique does not provide a mathematical tool that is useful in the quantization process.

2.2.2 Context-dependent colour mapping

Even though the main research in quantization is to find a good palette, the mapping of the original image to this palette is also important. As it was stated before, the current quantization methods are not able to account for colour context. Fortunately it is possible to overcome this limitation by the use of the context-dependent colour mapping — *dithering*. This technique is able to reduce the artificial contouring.

The nature of this artifact is explained by Figure 2.1. The colours c_i and c_j belong to two different Voronoi regions. The colour mapping (2.1) assigns corresponding




Figure 2.1: The nature of artificial contouring in quantization. The colour mapping increases the original contrast $\|c_i - c_j\|$ to $\|\bar{c}_i - \bar{c}_j\|$.

pixels to \bar{c}_i and \bar{c}_j . Thus in the displayed image the original contrast $\|c_i - c_j\|$ is increased to $\|\bar{c}_i - \bar{c}_j\|$.

Dithering methods are based on *the spatial integration* that our eyes perform. If we view a very small area from a sufficiently large distance, our eyes average the detail. An image region with a combination of pixels of different colours may appear to be uniformly shaded.

Most dithering algorithms are developed for tone reproduction in printing [Uli37]. The error diffusion techniques can be easily adapted to quantization colour mapping. The Floyd and Steinberg approach [FS75] is one of the most commonly used algorithms. The quantization error $\|c - q(c)\|$ is spread over the weighted neighbourhood (Figure 2.2). Hence the context dependent colour mapping operation for pixel $c_{(x,y)}$ can be defined by the equation:

$$q_{(x,y)}(c) = q(c + \epsilon_{(x,y)}) \quad (2.3)$$

where $\epsilon_{(x,y)}$ is the accumulated error in the (x, y) neighbourhood; $q(c)$ is the previously defined context-independent mapping (2.1).

The shortcomings of the Floyd-Steinberg approach appear as correlated artifacts, directional hysteresis etc. A detailed discussion of these artifacts and the possible solutions are discussed by Ulichney (see [Uli87], Chapter 8)




Figure 2.2: Error propagation in the Floyd-Steinberg algorithm

In the case of fixed palettes with uniformly distributed colours other dithering algorithms may be applicable. In particular it is advantageous to use less computationally expensive ordered dispersed dithering ([Uli87], Chapter 6).

The application of context-dependent colour mapping is demonstrated by Figure 2.3. The test image “Lenna” is quantized using the same palette with 30 uniformly distributed colours. The use of dithering resulted in improvement of the perceived colour approximation and reduction of artificial contouring. Unfortunately distortion measures that are able to account for dithering effects have yet to be developed.

2.3 Approaches to colour quantization

As it is seen from the previous sections colour quantization is a hard problem. The lack of the objective criteria of quantization accuracy makes it difficult to formulate the task analytically. Hence, the thesis presents quantization techniques that are highly heuristic and approximate. This section outlines the early approaches to the colour reduction problem.

Unlike multi-dimensional vector quantization, quantization on one parameter is a well studied problem. An optimal partitioning can be found by dynamic programming [Bru65]. The complexity of these methods for N -level input is $O(N^2K)$. The early colour quantization methods [Hec80] employed these techniques to quantize each colour channel independently. Such an approach is not able to account for the multi-dimensional nature of colour.

Stevens et. al. [SLP83] coded three-dimensional colour information by a one-

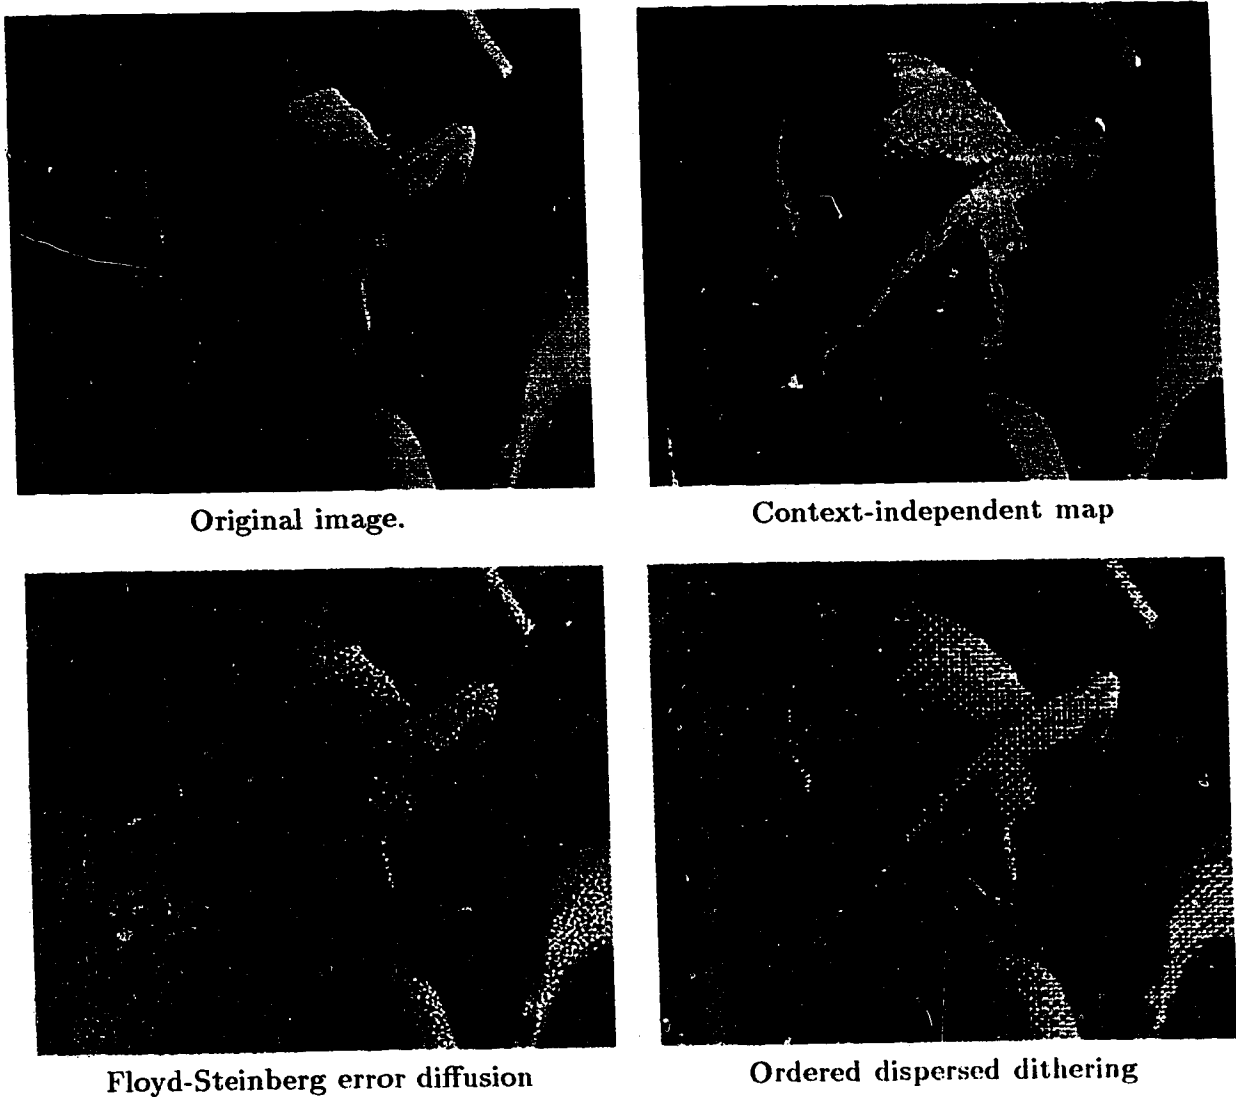


Figure 2.3: Application of dithering to quantization. The context-independent colour mapping significantly distorts the original image. The colour mapping with dithering results in reduction of artificial banding and improves the perceived colour approximation. However, dithering may introduce directional artifacts.




Figure 2.4: Peano scan encoding of the colour space preserves neighbourhood relationship between colours. Colours c_j and c_k are close in the colour space and along the space filling path. However the inequality $\|c_i - c_j\| < \|c_i - c_k\|$ is true in the colour space but false along the Peano scan.

dimensional parameter. This parameter represents the distance from the colour to the origin along the Peano scan. The quantization is performed on this distance. The method relies on the fact that a chosen fractal curve is able to preserve neighbourhood relationship between colours. That is, two colours c_j and c_k that are close along the Peano scan are also close in the colour space. (Figure 2.4). However the inverse statement is not valid. The colours c_j and c_i are far apart along the space filling curve, but are close in the colour space. Thus the relative distances are not preserved by this encoding. The inequality $\|c_i - c_j\| < \|c_i - c_k\|$ is true in the colour space but false along the Peano scan. Thus the resulting one-dimensional quantization is optimal for the chosen encoding but fails in the multi-dimensional space.

The early quantization algorithms showed that methods of one-dimensional quantization cannot be easily adapted to colour reduction problem. The multi-dimensional nature of colour is important to preserve. The complexity of the optimal partitioning problem forces us to find an approximate solution.

The *popularity algorithm* [Hec80] presents a somewhat simplistic approach to the problem in the 3-dimensional colour space. The assumption was made that a good

colour map should contain the most frequent image colours. The histogram of the image pixels is sorted using a selection sort until K most “popular” colours are selected. The computation complexity of this approach is $O(NK)$ where N is the number of colours in the histogram. The algorithm may work for some images but fails for those with uneven colour distribution. The method neglects colours in sparse regions of the space (see [Hec82]).

The rest of this thesis deals with quantization techniques that are designed to minimize one of the quantization errors defined in the previous section. We will distinguish two approaches in the algorithm design: *pre-* or *post-clustering* [Dek94]. The pre-clustering methods partition the space into regions relying on the pre-computed statistics of the colour distribution. The palette is chosen as centers of the generated regions. Post-clustering algorithms find representative colours first. These algorithms start with some initial approximation of the palette. This palette is iteratively improved based on multiple sample sets of the image pixels. The following chapters analyze merits of both approaches. The new local K-means algorithm is a post-clustering scheme. This algorithm is meant to overcome some limitations of the previous methods that follow this approach.

2.4 Algorithm evaluation methodology

This research compares a new technique with implementations of the quantization algorithms found in public domain image processing software: median-cut [Pos91], variance-based [Tho90], octree [Cri92], and SOM [Dek94]. These implementations work in RGB colour space. For a fair comparison we also used RGB space. Note, that even though quantization in perception-based spaces can give a better visual result, it does not change the relative correspondence of numerical values of quantization accuracy. Therefore, algorithms that produce small distortions in RGB space are expected to perform as well in Lu^*v^* or HSV spaces.

2.4.1 Visualization of quantization errors


For our quantization tests we have chosen a set of 24-bit images that represent various image sources: scanned photographs, computer rendered scenes, and digitized works of art.

The chosen test images are meant to represent the two discussed quantization artifacts (Figure 2.5). The image “Kiss” has a nonuniform distribution of colours. There are few pixels of red, green and blue in this mainly golden painting. Such colour diversity makes the image hard to quantize to a small palette. The relatively “rare” colours are often missing in the generated palettes. Thus the image “Kiss” demonstrates the loss of colour information and contrast.

The “Pool balls” computer rendering has large uniformly shaded areas. The artificial contouring is apparent even when this image is quantized to a 256 colour palette.

Due to the limited printing technology the quantization results are rather difficult to evaluate using hard copies of the test images. The artifacts may be visible on the computer monitor but indistinguishable in print. We have chosen to use the quantization error visualization technique proposed by Fiume et. al. in [FO89]. In this thesis gray scale images represent the difference between the original and the quantized colour images. The intensity of each pixel is computed as the Euclidean distance of the colour distortion: $\|c_{(x,y)} - q(c_{(x,y)})\|$. Therefore the brightest areas in these images correspond to the highest quantization errors.

It is clear that such a visualization technique is rather limited as it does not provide information about colour contrast. However, this methodology is suitable for the discussion of error distribution in the quantized images.



A digitized painting "Kiss"



Computer synthesized image "Pool balls"

Figure 2.5: The test images.

Chapter 3

Pre-clustering

Pre-clustering approaches are based on the formulation of the quantization problem as an optimal space partitioning, where the colour space is split into a set of regions. The goal of this division is to classify similar colours into the same cluster. Wu in [Wu92a] demonstrated that this objective follows from the definition of the optimal partitioning. Once a suitable set of clusters is generated the colour map is chosen as centroids of the resulting regions.

This chapter describes the pre-clustering algorithms used in computer graphics. These methods differ in the heuristic strategies of the space partitioning. We will analyze the influence of the different approaches on the quality of the quantized images.

3.1 Median cut

The objective of the median cut algorithm [Hec82] is to partition a colour space into clusters with equal number of pixels. Thus each entry of the synthesized colour map will represent equal number of pixels of the original image. The colour space is repeatedly subdivided into a set of rectangular boxes by planes parallel to the space axis.

Prior to the partitioning step the histogram of image colours is built. In order to be efficient the number of distinct colours is reduced. This is achieved by colour




Figure 3.1: The first 3 steps of the median cut algorithm. This technique attempts to partition the space into colour regions with equal number of pixels.


resolution reduction. Heckbert suggests that 5 bits per colour channel is sufficient in most cases.

The algorithm starts with one box that encloses colours of all pixels in the input image. The maximum and minimum values along each axis are found. The colours are sorted along the axis with the longest edge. The box is split at the median point by a plane perpendicular to this edge

Thus the number of pixels in the resulted sub-boxes is approximately equal (Figure 3.1). The operation is recursively applied to the new partition. The process stops when the needed number of clusters is achieved. The palette is generated from the centroids of colours in each box.

The median cut approach is able to produce good colour maps for images with approximately equal distribution of colours in the space. However, artificial contouring appears in the smoothly shaded areas of the image “Pool balls”. Unfortunately, the algorithm approximates more frequent colours at the expense of the distinct rare colours. Thus images with uneven colour distribution may look significantly distorted. (Figure 3.2).

Joy et. al. [JX93] suggested that the primary weakness of the median-cut method is the decision to split the most popular box. The flaw of this approach is that the box



“Kiss”



“Pool balls”

Figure 3.2: Quantization errors: median cut method. The more popular colours are well represented at the expense of other less frequent colours. Thus the image “Kiss” exhibits strong colour shift. The artificial banding is apparent in the image “Pool balls”.





Figure 3.3: The first 3 steps of the variance based algorithm. The method splits the box with the largest colour disparity. The splitting plane is chosen to minimize variance of colours in the resulting regions.

with the largest pixel count may have little colour disparity (small variance). Leaving less popular but larger boxes unsplit causes higher quantization errors. Moreover the objective of the space partitioning — each palette entry represents equal number of pixels — does not allow the median cut to take into account less frequent colours.

3.2 Variance-based algorithm.

Wan et. al. [WPW90] attempted to improve the performance of the pre-clustering scheme and proposed a new objective for space partitioning. Optimal quantization is possible if the variance of pixel values is minimal in each resulting cluster.

The variance based method [WPW90] follows a scheme similar to the median cut. At each step the box with the largest weighted variance of colours is selected. The weight of the box is the number of pixels it encloses. To choose the partition plane, the distribution of colours along each of three axis is computed. The optimal threshold is obtained for each of the projected distributions. The weighted sum of projected variances is computed for the three pairs of the possible partitions. The partition plane is chosen to be perpendicular to the axis with the smallest sum and to pass through the optimal threshold.



"Kiss"



"Pool balls"

Figure 3.4: Quantization errors: variance based method. The algorithm attempts to minimize variance within each colour cluster. This approach leads to more even colour approximation than the median cut. The colour shift and banding in the test images are less evident.

The subdivision continues until the needed number of boxes is created. As before, the mean colour of the box becomes an entry in the palette. Figure 3.3 demonstrates that the variance based algorithm is able to split the boxes in such a way that colours are well grouped. Thus the colour disparity within each cluster is small. Unlike the median cut partitioning, the variance based method creates boxes of different sizes. The close colours tend to end up in the same box. This results in the better approximation of less popular colours. Therefore, the colour approximation is more uniform than that of the median cut algorithm. The images of quantization errors (Figure 3.4) have less bright areas than the corresponding median cut images (Figure 3.2).

3.3 Center cut

The center cut algorithm is a simplified approach to variance minimization within colour regions. Joy et. al. [JX93] assumed that the box with the longest edge has the highest variance. The proposed center cut algorithm splits this box in half with the partition plane perpendicular to the longest edge. This center cut method produces palettes similar to the variance based method significantly faster.

3.4 Octree

Heckbert [Hec82] points out that the structure formed by recursive partitioning of the presented algorithms is nearly identical to k-d trees. The octree algorithm [GP88], [CFM93], [Cri92] explores this idea further.

The entire colour space is treated as a hierarchy of octants. Each colour of the input image is placed into the leaf of the constructed tree. The “pruning” starts with the longest “branch”. Neighbouring leaves are recursively substituted by their parents. The process stops when the number of leaves is equal to the required size of the palette. Thus the leaves become the entries of the colour map (Figure 3.5).

The octree algorithm is similar to the previous methods as each leaf is a box in the





Figure 3.5: One level of the octree partitioning. Similar colours may end up in different clusters.

colour space. We classify this method as a pre-clustering approach since it partitions the space prior to the palette selection.

The representation of the colour space by an octree allows memory efficient implementation of the quantizer. The colour resolution reduction and histogramming of image pixels is avoided. Hence, the octree method is significantly faster than the other techniques. The distinct colours are likely to be in different octree leaves. Therefore, the quantization results are somewhat similar to the variance based method (Figure 3.6, “Kiss”). However, the algorithm does not guarantee that relatively close colours are placed into the same box. The octree quantization may introduce strong banding artifacts. In the merging process the volume of the box grows by a factor of eight. This often creates imbalanced partitioning and leads to uneven colour gradations in the smooth areas of the quantized images (Figure 3.6, “Pool Balls”). Moreover the size of the palette for octree quantization is always a multiple of eight.

3.5 Principal axis splitting quantization.

The algorithms described above have a common flaw. The partitioning is performed by the orthogonal planes even though in general colour image data sets are not distributed orthogonally. It is advantageous to account for the *principal axis* of the



Kiss



Pool balls

Figure 3.6: Quantization errors: octree method.
The quantization results are similar to the variance based method. The strong banding is due to the fact that close colours may end up in the separate octree leaves.


colour distribution. This axis represents the direction of the maximum variance.

Wu [Wu92b] used a cutting half-plane normal to the principal axis of a data set. The cluster is thus split in the locally optimal manner. The position of the plane is chosen to minimize the quantization distortion. The inertial cut method is the only implementation of the principal axis partitioning technique available to the author. This method splits the region in the center of mass. This approach allows us to reduce the computation cost. The inertial cut is able to generate a high quality palette for the test image "Pool balls". In particular the artificial banding is practically invisible in the quantized image. Unfortunately the algorithm fails to approximate rare colours in the image "Kiss". Figure 3.7 presents quantization results for the inertial cut algorithm.


Even though the principal axis approach improves the colour reduction it has a serious limitation. The intermediate clusters are bipartitioned one at a time independently from each other. As a result the quantization process is not able to take into account interrelationships between neighbouring regions.

Wu [Wu92a] studied the principal multilevel quantization algorithm. The motivation of this scheme relies on the fact that the colour distribution in the natural scene is not isotropic in the colour space. The values are more spread out in the luminance direction (L component of the CIE Lu^*v^*). It was observed that the first partitions created by the previous algorithm had almost parallel sides. The objective was to optimize these first cuts by simultaneous planes perpendicular to the principal axis. The position of these planes is determined by a dynamic programming technique similar to the one used for one-dimensional quantization (Figure 3.8).

The experiments in [Wu92a] showed the advantage of the multilevel quantization. The partitioning was better adapted to the statistics of the input image. Wu claims that *"on average, the mean-square quantization error of the new algorithm is five times smaller than that of the traditional algorithms"*. The improvement in the colour approximation comes at the high computation cost: *quantization takes "less than three minutes on a Personal IRIS workstation"* [Wu92a].



Kiss



Pool balls

Figure 3.7: Quantization errors: inertial cut method. The algorithm is able to avoid artificial banding in the image "Pool balls". Even though the inertial cut generates low quantization distortion per pixel the distinct small regions of the image "Kiss" are poorly approximated. (bright areas of the error images: grass, flowers).




Figure 3.8: Principal multilevel space partitioning. The first step splits the space by multiple planes perpendicular to the principal axis. The simultaneous split allows optimization of the multiple cuts. The second step partitions the resulted region with the largest variance.

3.6 Wrapup

A survey of pre-clustering algorithms was presented in this chapter. These techniques subdivide the colour space into clusters according to the distribution of colours in the input image. The centroids of the resulting regions form the quantization palette. The objective of this space partitioning is to place similar colours into the same cluster. Taking into account that the quantization problem is NP complete [Wu92a], the described techniques offer an approximate solution. These methods share a similar recursive partitioning scheme. The algorithms differ in the splitting criteria and the direction of the cut.

These pre-clustering quantization methods are popular in computer graphics applications ([Pos91], [Tho90], [Cri92]). They are able to generate reasonable colour maps at a moderate computation cost. However this ease of computation often compromises the quantization accuracy. In order to be efficient the pre-clustering algorithms choose to reduce the colour resolution of the input image.

The minimization process of these recursive schemes is tied to the created regions of the space. However the closest palette entry for some colours may be different




Figure 3.9: Voronoi regions of the median cut partitioning. Colours depicted by gray on this diagram will be mapped to a different palette entry than the center of their partitioning region.

from the center of their cluster. Hence colours belonging to the same cluster in the space partitioning may be mapped to different palette entries in the quantized image (Figure 3.9). Thus the space partitioning may be optimal in respect to the generated clusters and produce high quantization errors.

Chapter 4

Post-clustering algorithms

This chapter is devoted to the post-clustering quantization approaches. These techniques offer a direct solution of the vector quantization problem. The algorithms that follow this approach start with some initial palette and iteratively improve it to minimize the quantization error. Unlike the pre-clustering schemes, these techniques do not require computation of various statistical parameters of the image. The adaptation process attempts to approximate the density function of the colours using multiple samples of the input. These algorithms have been applied to statistical analysis, data coding, signal processing and pattern recognition [LBG80], [Gra84], [Fri93b], [MG93], [KKL90]. Until recently these schemes were considered to be too computationally expensive for colour quantization.

4.1 K-means algorithm

K-means algorithm [LBG80] is a post-clustering technique that is widely used in image coding and pattern recognition. A sequence of iterations starts with some initial set $\bar{C}^{(0)}$. At each iteration, t , all data points $c \in C$ are assigned to one of the clusters $S_k^{(t)}$. The cluster membership is defined by the closest center from the set $\bar{C}^{(t)}$. The centroid of all points $c \in S_k^{(t)}$ becomes a new center of the cluster $\bar{c}_k^{(t+1)}$:

$$\bar{c}_j^{(t+1)} = \frac{1}{l} \sum_{i=1}^l (c_i | c_i \in S_j^{(t)}). \quad (4.1)$$

The algorithm is known to converge to a local minimum.




Figure 4.1: K-means adaptation scheme. All data points are assigned to clusters defined by the closest center (black dots). The centroids of the resulted clusters become the new centers (gray dots).

The K-means algorithm was used to quantize images in [WPW90]. For the test images it produced smaller average errors E than the median cut and variance based pre-clustering algorithms. Unfortunately, high cost of computation makes K-means impractical for image quantization.

4.2 Kohonen self-organizing maps.

A self-organizing map (SOM) was introduced by Kohonen [KKL90] as a solution to a general vector quantization problem. The SOM is a neural network that imposes a one or two-dimensional topological structure over a set of clusters in a higher dimensional space. Dekker in [Dek94] studied the use of the one-dimensional self-organizing map for image quantization. The initial palette is set to equally spaced gray scale values. The input colours are obtained by multiple sampling of the image with large step sizes. The closest colour $\bar{c}_k^{(t)}$ of the palette is adjusted to better comply with the input $c^{(t)}$. The adaptation process is controlled by the adaptation parameter $0 < \alpha_t < 1$. This parameter is exponentially decreasing with time thus




Figure 4.2: SOM adaptation scheme

the iteration process converges to a stable set of clusters.


The network is considered to be *elastic*. That is the $2r$ topological neighbours are updated together with the center $\bar{c}_k^{(t)}$. The parameter r is the radius of elasticity that decreases with time. The elasticity coefficient $\rho_{(t,j)}$ ensures that only entries in the r -neighbourhood are updated. Since the updated neighbourhood often overlap the values of \bar{c}_k tend to be smoothed.

To summarize the SOM adaptation process is presented on Figure 4.2 and defined mathematically by the following equation:

$$\bar{c}_j^{(t+1)} = \bar{c}_j^{(t)} + \alpha_t \rho_{(t,j)} (c^{(t)} - \bar{c}_j^{(t)}), \quad (4.2)$$

In order to ensure a fair representation of colour regions by the palette \bar{C} Desieno (see [HN90] p. 69) proposed the use of a special bias value. The bias factor increases for less frequently chosen vectors. Thus a colour that was chosen many times before has a lower probability to be chosen later.

Figure 4.3 presents the results of the SOM quantization of the image “Pool balls”. The quantized image does not contain strong contouring artifacts. The technique favorably compares with octree, median cut and variance based pre-clustering schemes. Unfortunately the SOM quantization is significantly slower than other algorithms (SOM requires 48 sec. verses 2-3 sec. for the tested pre-clustering methods). Dekker proposed to use only a part of the image as an input data to generate the palette.



Fully adapted network for image "Pool balls"



Quantization errors for image "Pool balls"

Figure 4.3: SOM quantization test.
The colours of the original image are well represented by the selected palette. The artificial contouring is avoided.

```

Initialize the temperature:  $T = T_0$ 
while ( $T > T_1$ )
   $accepts = rejects = 0$ 
  while ( not at equilibrium : ( $accepts < Max_{accepts}$ )( $rejects < Max_{rejects}$ )
    perturb one of the palette entries
    compute the change in the global quantization error  $\Delta E$ 
    if ( $\Delta E < 0$ )
      accept the new palette
       $accepts = accepts + 1$ 
    else
      choose  $random \in [0, 1]$ 
      if ( $random < e^{-\Delta E/T}$ )
        accept the new palette
         $accepts = accepts + 1$ 
      else
         $rejects = rejects + 1$ 
  Reduce the temperature:  $T = Tk$ 

```

Figure 4.4: Simulated annealing quantization

This approach speeds up the palette selection but reduces quantization accuracy.

4.3 Simulated annealing

A number of probabilistic techniques were developed to solve large scale optimization problems. One of these techniques is *simulated annealing*. This method is based on the analogy from the statistical thermodynamics. Simulating annealing mimics the process of crystallization of a liquid under low temperatures. The cooling schedule ensures that crystallization occurs at the lowest possible energy level. This is achieved if the temperature decreases slowly in a sequence of discrete steps. Thus the equilibrium is reached at each temperature level.

Fiume et. al. [FO89] applied simulated annealing to colour image quantization. In this case, the energy to minimize is the global quantization error. The initial temperature is computed from the initial approximation of the palette. The overall scheme of the annealing algorithm is presented on Figure 4.4

This technique was applied to improve colour maps generated by one of the pre-clustering methods [FO89]. The distributed version of the annealing algorithm was used. In the course of experiments it was found that this probabilistic scheme was able to reduce the global quantization error. However, the improvement was possible only after thousands of iterations (several hours on the Sun 3/280). Besides the slow convergence rate it was noted that “it is difficult to devise scientific ways of determining good values for parameters ” in the cooling schedule. The authors found acceptable values by trial and error.

4.4 Motivation for the future research

The post-clustering methods share a common iterative scheme. That is the colour map is iteratively adapted to minimize the global quantization error. The adaptation terminates when the needed approximation accuracy is reached. The author believes that the iterative nature of post-clustering techniques might be appealing in many computer graphics applications. Computation time or quantization accuracy could be adapted depending on the application requirements.

Unlike many pre-clustering methods, the post-clustering algorithms tend to minimize the global error by simultaneous minimization of all partitions. Thus, the quantization error is distributed more evenly for all colours of the image ([Wu92a], [FO89]).

As it was noted before the existing post-clustering methods appear to be too slow for practical use in computer graphics. Hence, the objective of this research is to derive a fast algorithm that follows the post-clustering scheme.

Further chapters describe the original results presented by the author in [Ver95] and [VB95].

Chapter 5

Fast nearest neighbour search.

The performance of a quantization method greatly relies on the speed of the nearest neighbour search. This search is the basis of the colour mapping operation. Moreover, the described post-clustering techniques use the nearest neighbour to determine the optimal palette.

In order to speed up the search Freidman et. al. proposed the use of k-d trees [FBF77]. It was proven that the complexity of this search is $O(M \log K)$. In his software Poskanzer implements the colour mapping using a hash table (see [Pos91]). Heckbert in [Hec82] described the *locally sorted search*. The colour space is divided into a set of cubical cells that contain a list of palette entries. The representative colour is inserted into the list if it is the nearest neighbour of one of the points in the cell. These lists are sorted by the distance from the cell. The sorting allows terminating the search for the closest colour before the entire list is examined. On average, the proposed technique is 23 times faster than exhaustive search.

Unfortunately many algorithms developed for fixed colour maps cannot be used in the framework of iterative post-clustering procedures. Positions of representative colours \bar{c}_j are constantly changing, therefore a k-d tree or a hash table must be recomputed after every iteration.

Hodgson [Hod88] proposed several speed up techniques for the minimum distance classifier: the partial sum test, sorting the cluster centers, calculation of the nearest neighbour for each center. In this chapter we discuss applications of these

optimization approaches to colour quantization. In addition the author proposes to compute classification in respect to the L_α norm. This inexpensive approximation to the Euclidean metric is presented in the following section.

5.1 L_α -norm

Previously we have defined the colour mapping and minimal distance classification in terms of the Euclidean metric. Unfortunately this metric is computationally expensive. Many implementations of colour quantization algorithms substitute the Euclidean L_2 norm by the less expensive L_1 norm ([Dek94], [Pos91], [Tho90]). However the nearest colour determined by L_1 norm may not be the nearest colour in L_2 norm.

Chaudhuri et. al. [CCW92] proposed the L_α norm as an approximation of the Euclidean metric. For a vector $x \in R^n$ the L_α norm is defined as a combination of the L_1 and L_∞ metrics:

$$\begin{aligned} \|x\|_\alpha &= (1 - \alpha)\|x\|_1 + \alpha\|x\|_\infty \\ &= (1 - \alpha) \sum_{i=1}^n |x_i| + \alpha \max_i |x_i|. \end{aligned} \quad (5.1)$$

The choice of the $\alpha = 1/2$ simplifies the norm calculations. In this case $\alpha = 1 - \alpha$, thus the multiplication can be avoided in the search of the nearest colour. The author found that the application of the $L_{\alpha=1/2}$ norm significantly speeds up the colour mapping (Table 5.1). Moreover the resulting misclassifications do not noticeably influence the quality of the output image.

5.2 Optimization techniques

This section describes search optimization techniques proposed by Hodgson in [Hod88]. These methods are possible to apply in the framework of the post-clustering algorithms when the palette entries are changing with iterations.

Table 5.1: Application of the $L_{\alpha=1/2}$ to the colour mapping. The $L_{\alpha=1/2}$ takes only 24% of the time to compute the L_2 norm. However the $L_{\alpha=1/2}$ based colour mapping is similar to the Euclidean mapping. The less expensive L_1 metric introduces high average distortion per pixel \mathcal{E} due to a large number of misclassifications (11%).

Norm	Time	\mathcal{E}	Wrong neighbour
L_1	11.9 sec.	5.56	11%
L_2	59.9 sec.	5.46	
$L_{\alpha=1/2}$	14.7 sec.	5.47	4%

5.2.1 Calculation of the partial sum

It is often the case that a partial calculation of the distance is sufficient to rule out the current palette entry as the closest colour to the input c . The norm calculation is abandoned if the current partial sum Σ_p exceeds the current minimal distance Σ_{\min} .

Consider the situation on Figure 5.1. Let the current minimal distance be $\Sigma_{\min} = \|c - \bar{c}_1\|$. The partial sum — the distance in horizontal coordinate — for \bar{c}_3 exceeds the Σ_{\min} . Therefore the complete calculation of the norm is not necessary for this colour. The colour \bar{c}_4 is the closest palette entry to the current input.

5.2.2 Sorting on one coordinate

Sorting eliminates a number of centers from the search for the closest palette entry.

Suppose that the palette colours are sorted according to their projections on the horizontal axis (Figure 5.2). The projection of \bar{c}_4 is the closest to the projection of the input c on the chosen axis, thus the search starts with \bar{c}_4 . The current minimal distance is $\Sigma_{\min} = \|c - \bar{c}_4\|$. The search needs to examine only the entries with projections within Σ_{\min} from the input. These colours are visited in alternating order from the left, and from the right of the input — $\bar{c}_3, \bar{c}_5, \bar{c}_2, \bar{c}_1$. Since $\|\bar{c}_3 - c\| < \Sigma_{\min}$ the \bar{c}_3 is the new closest entry. The horizontal projection \bar{c}_2 is outside of the current minimal distance radius. Therefore \bar{c}_2 and all the other entries to the left of it are




Figure 5.1: The application of the partial sum criteria. If the $\Sigma_{\min} = \|c - \bar{c}_1\|$ is the current minimal distance, then the complete calculation of the sum is not necessary for \bar{c}_3 .

eliminated from the search.

5.2.3 Nearest neighbour distance (NND)

The nearest neighbour criteria is another optimization technique that allows us to terminate the search without examining the entire palette.

The nearest neighbour is found for all the palette entries, and half of this distance is stored. Let us suppose that \bar{c}_3 is currently the closest entry to the input c (Figure 5.3). The entry \bar{c}_2 is the nearest neighbour of the entry \bar{c}_3 . Since the input c is in the sphere of radius $\frac{1}{2}\|\bar{c}_3 - \bar{c}_2\|$, then the search can terminate. The other palette entries cannot be closer to c than \bar{c}_3 .

5.2.4 Experiments in colour mapping

The effectiveness of the described search optimization techniques is studied in application to colour mapping. A 512x400 image was quantized to 16 and 256 colour palettes. Computation times can be found in Table 5.2.




Figure 5.2: Sorting by one coordinate. The entries whose projections are outside of the current minimal distance are eliminated from the search. Thus \bar{c}_2 and \bar{c}_1 are not considered.




Figure 5.3: The application of the nearest neighbour criteria. The \bar{c}_2 is the nearest neighbour of \bar{c}_3 . The $\frac{1}{2}\|\bar{c}_3 - \bar{c}_2\|$ exceeds the distance between \bar{c}_3 and the input. Thus other palette entries cannot be closer to c than \bar{c}_3 . The search terminates and \bar{c}_3 is the closest palette entry to c .

Table 5.2: Performance of the colour mapping optimization techniques. The optimized colour mapping is comparable to the $k - d$ tree search.

Algorithm	Execution time	
	16 colours	256 colours
Direct search	1.74	25.51
\sum_p	1.45	19.44
\sum_p and sorting	0.89	2.93
\sum_p , sorting and NND	0.76	2.76
k-d trees	1.41	2.65

The implementation of all optimization criteria resulted in 10 fold speed up of the colour mapping operation. The performance of the $k - d$ tree search is given as a reference. Even though this search has a logarithmic complexity the advantages of this algorithm are superseded by the required overhead.

Thus according to the experiments the performance of our colour mapping algorithm is comparable to k-d tree method.

Chapter 6

Local K-means for colour quantization.

The advantages of the post-clustering quantization techniques motivated the development of the local K-means algorithm. The main goal of this research was to investigate optimization techniques that would result in an efficient post-clustering approach to colour quantization.

This chapter describes the proposed local K-means algorithm. This technique is applied to a number of test images. The quantization examples and error statistics are presented to illustrate both the speed and accuracy of the proposed method.

6.1 The local K-means algorithm

The local K-means algorithm is an iterative adaptation process, similar to the other post-clustering methods. The algorithm starts with the initial palette. Multiple samples of image pixels are used to adapt the colour approximation. This section describes three aspects of the local K-means algorithm: the adaptation process, the choice of the initial palette, the image sampling technique.

6.1.1 The adaptation procedure.

Let $c^{(t)}$ be an input colour on the iteration step t . Suppose that the palette entry $\bar{c}_k^{(t)}$ is the closest to the input $c^{(t)}$. The local K-means algorithm adapts the colour $\bar{c}_k^{(t)}$ to better approximate the input. The parameter α_t guarantees the convergence of the iteration process. Thus the adaptation procedure of the local K-means algorithm can be expressed by the following equation:

$$\bar{c}_j^{(t+1)} = \begin{cases} \bar{c}_j^{(t)} + \alpha_t(c^{(t)} - \bar{c}_j^{(t)}) & j = k; \\ \bar{c}_j^{(t)} & \textit{otherwise} \end{cases} \quad (6.1)$$

The local K-means is similar to gradient quantization techniques used in gray scale image coding [Mat92], [MC92]. Also, the local K-means algorithm can be considered a special case of a self-organizing map. Unlike the Kohonen network, the adaptation step of the LKM process updates only the closest colour. This approach eliminates the need to maintain the topological map — the neural network of the palette entries. The modification results in a speed up of the iteration process.

The local updates of the local K-means procedure are advantageous for quantization of images with uneven colour distribution. Suppose the colour $\bar{c}_j^{(t)}$ is the only representative for a small group of distinct image colours. It is likely that its topological neighbours $\bar{c}_{k\pm 1}^{(t)}$ are a significant distance from the current input $c^{(t)}$. The SOM algorithm will erroneously update these distant palette entries. Thus the overall quantization accuracy may decrease.

6.1.2 The adaptation parameter

The adaptation parameter $0 < \alpha_t < 1$ (6.1) ensures the convergence of the iteration process. Kohonen in [Koh91] suggested that the parameter α_t should satisfy the following conditions:

$$\sum_{t=0}^{\infty} \alpha_t = \infty, \quad \sum_{t=0}^{\infty} \alpha_t^2 < \infty \quad (6.2)$$

We have chosen to follow the approach proposed by Dekker in [Dek94]. He suggests the use of $\alpha_t = e^{-0.03t}$. Even though this adaptation parameter does not satisfy the above conditions, it performs well for colour quantization. The number of experiments in [Dek94] show that the proposed α_t reduces a number of iterations needed to find a good palette.

It is important to point out that a different choice of α_t may result in a different final palette. We found that it is advantageous to choose the adaptation parameter to be dependent on the number of pixels M in the image. The current implementation of the local K-means algorithm uses an α_t of the form:

$$\alpha_t = \alpha^t : \quad \alpha = \frac{M - \delta}{M} \quad (6.3)$$

where δ is a number between 15 and 25. The speed of the adaptation will increase with larger δ . A “very fast” adaptation may lead to the reduction of quantization accuracy. Therefore the user can tailor the quantization process to the specifics of the current application.

6.1.3 The choice of the initial palette.

The final result of the adaptation process is greatly influenced by the selection of the initial palette. A number of techniques have been developed to lessen this dependence [KKL90], [PBT93], [Fri93b]. Unfortunately the proposed methods are computationally expensive and may not be suitable for fast colour quantization.

All the entries of the initial approximation $\overline{C}^{(0)}$ should be chosen inside the possible range of the input values. The vector quantization literature suggests several approaches. If the range of the input values is known the initial set $\overline{C}^{(0)}$ can be a sequence of equally spaced values within this range. Thus, Dekker in [Dek94] initializes the palette to gray scale values. Linde et. al. [LBG80] proposed to generate the “*initial guess by splitting*”. The set $\overline{C}^{(0)}$ of K values is chosen incrementally by recursive quantizing the data to $K/2$ clusters and perturbing their centers. We have studied a computationally inexpensive procedure similar to the the K-means

approach. The first K colours of the random image sample are inserted into the initial palette. Such a procedure guarantees that all the values of the $\overline{C}^{(0)}$ are inside the image gamut.

For an equal approximation of all colours in the image it is important that every distinct small colour cluster be represented by a separate palette entry. We found that this can be obtained if the colours of the initial palette are well separated. In the current implementation of the algorithm a new colour is added into the palette if its distance from the already inserted entries exceeds a specified threshold.

6.1.4 Input data.

The data for the adaptation procedure is a sequence of input sets. These sets are constructed by sampling the image in decreasing step sizes: 1009, 757, 499, 421, 307, 239, 197,... We have chosen these step sizes to be prime numbers, thus the input sets do not intersect too much. The iteration process stops when changes to the palette $\overline{C}^{(t)}$ over a complete image scan become small. In our experience the union of input sets does not include more than 10% of all image points.

Even though this algorithm examines only a portion of the input image, it is able to generate good approximating palettes. This result may be explained by the fact that colours of a typical image are clustered in the colour space. Therefore, it is enough to use a few colours from the cluster to approximate all its members. Since similar colours are often close to each other on the image surface, we hope that our input sets contain representatives for most clusters.

The local K-means algorithm allows the user to tailor the adaptation process to the image. For example, the user can specify critical areas of the image that need high approximation accuracy. The pixels from this critical area will be frequently included into the input sets. Thus a better approximation of colours in this area is expected.

Table 6.1: Quantization errors: image “Kiss”, 16 colours.

Method	Max	ϵ	\mathcal{E}	σ
Median cut	161	31.0	20.77	14.71
Variance based	146	25.4	17.63	11.57
Octree	133	26.4	19.41	13.65
Inertial Cut	133	26.2	17.74	11.49
Local K-means	106	20.8	26.65	9.30

6.2 Quantization results.

The local K-means algorithm is compared to other quantization techniques. We studied the approximation accuracy of the method in respect to two common quantization artifacts: the loss of colour information and the artificial banding.

6.2.1 Loss of colour information.

The reduction of the colour contrast and shift in the image hues are more apparent in quantization to a small number of colours (e.g. 16). These artifacts were examined by quantizing the digitized painting by Gustav Klimt “Kiss”. We have computed the measures of quantization distortion for all tested algorithms. (Table 6.1). The numerical values of the average distortion per colour ϵ and deviation of distortion per pixel σ are the smallest for the local K-means method. Therefore this algorithm is able to approximate all the colours equally well. This observation is supported by the quantization error images (Figure 6.1). The local K-means palette selection results in equal distribution of quantization errors. However, the uniform colour approximation leads to the higher average distortion per pixel \mathcal{E} .

Figure 6.2 demonstrates that local K-means procedure seem to reproduce the full chromatic range of the original image. High values of \mathcal{E} are apparent in the overall slight shift of image hues. Even though the quantized image produced by the median-cut method corresponds to a small \mathcal{E} it looks significantly distorted. Some fine details of the image have disappeared: blue flowers on the woman’s head,







Figure 6.1: The quantization error images for 16 colour palettes of image “Kiss”. The local K-means palette selection results in equal distribution of quantization errors. The uniform colour approximation leads to higher average distortion per pixel \mathcal{E} .




The original image



Median cut



Inertial cut



Local K-means

Figure 6.2: "Kiss": quantization to 16 colours.

The local K-means palette selection results in the smallest average distortion per colour ϵ and the highest average distortion per pixel \mathcal{E} . Such approximation better preserves the original image contrast. The smallest \mathcal{E} was produced by the inertial cut quantization. However some distinct rare colours are significantly distorted (e.g. green grass, yellow flowers).

Table 6.2: Quantization errors: image “Pool Balls”, 256 colours.

Method	Max	ϵ	\mathcal{E}	σ
Median cut	92	8.3	4.39	2.49
Variance based	45	6.6	4.27	1.72
Octree	61	6.4	2.07	2.16
Inertial Cut	43	6.0	1.33	2.05
Kohonen SOM	95	7.9	1.74	2.87
Local K-means	105	7.4	2.01	2.59


yellow spots on her dress, etc. The variance-based, octree and inertial cut algorithms were able to preserve most of these details, though the original colour contrast was greatly reduced. The subjective evaluation of quantization results shows that viewers are more sensitive to the change in the colour variation. Hence the local K-means quantization is often preferred over other techniques.

6.2.2 Artificial contouring.


The loss of colour information is not very apparent when images are quantized to large palettes. Unfortunately the images with large uniformly shaded areas are prone to artificial banding (or contouring).

A computer synthesized image “Pool balls” was quantized to 256 colours. All the tested algorithms were able to preserve the original colour contrast. The median cut, octree and variance based methods introduced significant banding. The local K-means was able to avoid this artifact. The smallest values of the quantization errors in Table 6.2 correspond to the inertial cut technique. Though the image produced by this method does not noticeably differ from the results of local K-means and SOM algorithms.

The statistical measures of quantization distortion in this case do not reflect the perceived results. The minimization of these parameters does not guarantee that the contouring will be avoided. Even though the statistical measures for octree and local K-means are similar, the later technique does not introduce the bending artifact



Octree



Local K-means

Figure 6.3: The quantization error images for 256 colour palettes of image "Pool Balls". The local K-means algorithm does not introduce noticeable artificial contouring. The statistical measures of quantization errors for local K-means and octree methods are similar. The banding is apparent for the octree technique.

(Figure 6.3). The context dependent measures of image distortion are needed for a more reliable control of image distortion.

6.2.3 Compromises

Our visual system is well trained to recognize human faces. We are able to distinguish small variations of colour and shape of eyes, lips, etc. Therefore accurate quantization of images with human faces is very important. Thus quantization to a small palette is especially challenging. The question arises: *“What is more important in minimization of the perceivable quantization distortion: accurate colours or smooth shading ?”* The portrait of the Hong Kong actress Anita Yuen (Figure 6.4) is chosen for this discussion. The quantization errors for palette with 32 colours are presented in Table 6.3.

The smallest average distortion per pixel \mathcal{E} was produced by the inertial cut quantization. The palette generated by local K-means algorithm introduced the smallest average distortion per colour ϵ and the smallest variation in quantization accuracy σ . The corresponding images are presented in Figure 6.5. The difference in colour approximation accuracy is visible on the lips of the actress (“red” for local K-means and “brown” for inertial cut quantization). However the penalty for good colour approximation is slightly increased artificial banding in smoothly shaded areas.

The author believes that for quantization to a small palette it is important to preserve the original colour gamut of the image and approximate all colours equally well. The artificial borders may be reduced by the error diffusion techniques. Figure 6.6 demonstrates that Floyd-Steinberg dithering eliminates the contouring artifact. However dithering is not able to compensate for poor colour approximation of the inertial cut method (the lips remain “brown”).

Thus we can conclude that small colour distortion is often favorable over apparent artificial contouring.





Figure 6.4: A digitized photograph of Hong Kong actress Anita Yuen Wing Yee


Table 6.3: Quantization errors: the Anita Yuen portrait, 32 colours

Method	Max	ϵ	\mathcal{E}	σ
Median cut	82	22.9	10.48	8.04
Variance based	66	18.4	8.96	5.64
Octree	58	14.7	9.41	5.62
Inertial Cut	79	15.3	8.31	5.40
Local K-means	73	12.5	10.75	4.20



Inertial cut quantization


The algorithm produces the smallest average distortion per pixel \mathcal{E} . However, the non-uniform colour approximation results in the noticeable colour shift (i.e. lips)




Local K-means

The algorithm produces the smallest average distortion per colour c . Unfortunately, the price for a good colour approximation is the increased artificial banding.

Figure 6.5: Quantization to 32 colours (the portrait of Anita Yuen)




The fragment of the original image



Inertial cut quantization

The dithering method is not able to compensate poor colour approximation (the lips remain 'brown').



Local K-means

The artificial banding is reduced by the error diffusion technique.

Figure 6.6: Quantization to 32 colours with Floyd - Steinberg error diffusion

Table 6.4: Execution time in seconds for colour map selection


Algorithm	“Kiss”		“Pool Balls”	
	612,096 pixels		195,330 pixels	
No. of colours:	16	256	16	256
Median-cut	4.64	4.93	1.2	1.57
Octree	2.04	3.67	0.65	0.90
Kohonen SOM		101.81		32.27
Local K-means	0.65	1.74	0.20	0.38

6.3 Performance optimization.


The design goal of the local K-means algorithm is to find a fast post-clustering quantizer. The adaptation process updates the closest palette entry to the current input. Thus, the search of the nearest neighbour is the bottleneck of this algorithm. We studied the application of the optimization techniques described in Chapter 5.

The tests demonstrated that palette sorting allows us to eliminate a large number of colours from the search. Even though the palette entries change, the adaptation process does not greatly rearrange the sorted order. It is likely that the need to resort decreases with time as the adaptation parameter α_t decreases. The current implementation of the local K-means method sorts the palette at the start of a new image sample. The size of the sample sets increase in time. Therefore the sorting is performed less frequently.

Overall the local K-means algorithm is able to select a colour map significantly faster than the other methods (Table 6.4). The dependence of the algorithm performance on the problem size is tested using the “Lenna” image. Figure 6.7(a) presents computation time required to quantize several scaled versions of this test image. Time for quantization to 16, 32, 64, 128, and 256 colour palettes are presented for the 512x400 version of the “Lenna” image by the graph in Figure 6.7(b). These graphs demonstrate that the execution time of the local K-means scheme is linear in respect to the image and palette size.



a) Dependence of the execution time on the size of the image



b) Dependence of the execution time on the output palette size

Figure 6.7: Complexity analysis

6.4 Wrapup

The quantization results show that the local K-means algorithm is a suitable approach to colour quantization. This research was focused on the tailoring of the post-clustering technique to the specifics of the quantization problem. The speed and accuracy of the algorithm were also studied.

The local K-means scheme is able to find good quality palettes. The advantages of the method are more apparent in the case of quantization to a small number of colours. The proposed technique attempts to equally approximate all colours of the image. This property is most important for images with uneven colour distribution.

The local K-means algorithm is significantly faster than the other tested methods. This performance was made by the application of different search optimization techniques.

Chapter 7

Multiple Image Quantization with Local K-means Algorithm.

Previous chapters of this thesis dealt with palette selection algorithms for the best display of one image. With the growing popularity of windowing systems the display of multiple images is required. Graphics applications in such systems must share display resources. In particular the colour map is shared between a number of simultaneously displayed images and windows' borders, icons, etc. This chapter explores the colour quantization problem in windowing systems. The local K-means algorithm can be adapted to account for previously allocated colours in the shared palette.

7.1 Colour quantization of multiple images

In order to approximate an image the previous quantization algorithms select a reduced colour set \bar{C} . This set can contain any element of the chosen colour space without any restriction. However simultaneous display of multiple images requires that the chosen palette \bar{C} be shared by the windowing system and the images. Therefore palette selection within the windowing system is a special instance of the quantization problem.

Let us assume that \bar{C}_S is the current shared palette with K_S colours. In order to display a new image I the palette \bar{C}_I with K_I colours is chosen. Unfortunately

quantization to a small number of colours may significantly distort the appearance of the displayed image. The colour approximation may be improved if the shared palette is also used. Thus we can assume that the image is quantized to K colours, where $K = |\overline{C_S} \cup \overline{C_I}|$. Since the sets $\overline{C_I}$ and $\overline{C_S}$ are chosen independently their intersection may not be empty, therefore in general $K \neq K_S + K_I$. Moreover, it is likely that the image specific palette $\overline{C_I}$ contains colours similar to those previously allocated in the shared palette. Thus the image approximation by the palette $\overline{C} = \overline{C_S} \cup \overline{C_I}$ is not very effective. This research attempts to answer the question: *How can a quantization algorithm account for the previously allocated colours ?*

We define the following “windows quantization problem”:

Complement a predefined set $\overline{C_S}$ of K_S colours with additional K_I elements such that the resulting set \overline{C} is optimal for quantization of image I .

The author believes that pre-clustering algorithms are not suitable for efficient palette selection within a windowing system. A recursive subdivision is not able to account for the allocated colours in the shared palette. The following section describes a modified local K-means algorithm that is appropriate for multiple image quantization.

7.2 Modified local K-means quantization

The goal of this research is to account for the shared colour map $\overline{C_S}$ in the local K-means adaptation process. The algorithm attempts to find the optimal approximating palette \overline{C} with K entries by adjusting “image specific” $K_I = K - K_S$ colours.

The K_S colours of the shared palette are inserted into the initial set $\overline{C^0}$. The initial values of the “image specific” palette entries are chosen by sampling the image. As it was previously discussed in section 6.1.3, the new entry is inserted into the initial colour set if it is a specified distance apart from other palette entries. Thus the palette entries are well separated.

The iteration t adapts only the “image specific” $\overline{C_I^{(t)}}$ subset of the current palette $\overline{C^{(t)}}$ to improve image approximation. The entries of the shared palette $\overline{C_S}$ are used

Table 7.1: Quantization errors for local K-means algorithm applied in windowing systems, image “Kiss”

Colour set	No. of colours	Max	ϵ	\mathcal{E}	σ
Local K-means					
$\overline{C_I}$	32	85	17.8	17.8	6.52
$\overline{C_I} \cup \overline{C_S}$	104	72	15.2	17.0	6.02
Windows local K-means					
C	104	72	14.0	15.1	5.52

in the adaptation process but remain unchanged. Thus for an input $c^{(t)}$ the modified local K-means procedure is expressed as follows:

$$\overline{c}_j^{(t+1)} = \begin{cases} \overline{c}_j^{(t)} + \alpha_t(c^{(t)} - \overline{c}_j^{(t)}) & j = k \wedge \overline{c}_j \in \overline{C_I^{(t)}}; \\ \overline{c}_j^{(t)} & \text{otherwise} \end{cases} \quad (7.1)$$

where α_t is the adaptation parameter described in section 6.1.2. The algorithm terminates when changes of the palette entries become small. All the colours in the new common palette will take equal part in the approximation of the original image, as if there was no restriction of the windowing system.

7.3 Experiments with quantization for windows systems

The experiments described in this section are meant to illustrate the use of the modified local K-means quantization in the windowing system. These tests are intended to simulate commonly used techniques in colour allocation and mapping (e.g. Mosaic for X Windows).

It was assumed that the windowing system is able to display $K = 256$ simultaneous colours and reserves for its own use $K_S = 8$ cells. A simulated windows application generates a small palette with 32 colours for each displayed image. These small palettes are added to the shared colour map of the windowing system, thus decreasing a number of free cells by 32.

In the course of the experiment three images were used. Table 7.1 presents the approximation errors for the third quantized image. At first the original local K-means algorithm was used. The image specific palettes are combined with the system palette into the global shared map, thus for the third test image \overline{C}_I contained $K_I = 32$ and the shared map \overline{C}_S — 72 entries. Even though the original local K-means selected the palette \overline{C}_I independently from the shared map, the colour approximation was improved when the test image is quantized to the palette $\overline{C}_I \cup \overline{C}_S$ (first rows in Table 7.1).

On the second step of this experiment the modified local K-means technique was applied to the same image sequence. 32 colour image palettes are chosen taking into account the current shared colour map. Table 7.1 demonstrates the improved colour approximation of the original image. Unfortunately the advantages of the modified local K-means procedure are not evident when the corresponding images are reproduced by the available printer.

7.4 Discussion

This chapter discusses the specifics of the windows quantization problem. The local K-means can be adapted to account for previously allocated colours. Therefore, the colour approximation may be improved. Moreover the entries of the global palette are well separated in the colour space, hence a large number of images can be approximated.

The tests of the modified local K-means technique were intended to imitate the colour allocation strategy of the current Windows application — The “image specific palettes” \overline{C}_I are of the same size. However the quantization accuracy is likely to increase if the shared palette is large. Thus to increase the number of simultaneously displayed images with the same quantization accuracy the size of \overline{C}_I should decrease. The size of the image specific palette sufficient for the specified quantization accuracy can be computed in a similar fashion as proposed by Fritzke in [Fri93b] and [Fri93a]. The new palette entries are created depending on the current quantization error.

Conclusion

This thesis presents the image quantization problem and discusses various approaches to effective colour reduction. This document includes a detailed survey of previous colour quantization techniques. The merits of these algorithms are analyzed with respect to the common quantization artifacts: the loss of colour information and artificial contouring. The computation efficiency is also considered.

Previously developed pre-clustering methods attempt to find an optimal space partitioning on the basis of the statistical distribution of colours in the image. These algorithms are highly heuristic but allow fast selection of a reasonable palette. Even though the iterative post-clustering methods generate more accurate approximations, they are considered computationally expensive for practical use in colour quantization.

The main contribution of this work is the development of a new efficient post-clustering algorithm known as local K-means. This method selects good quality colour maps significantly faster than previous pre-clustering schemes. This speed up was partly due to the implementation of nearest neighbour search optimization techniques. The advantages of the local K-means algorithm are more evident for quantization to a small number of colours. The selected palettes seemed to reproduce the chromatic range of the image and preserve the original colour variance.

The local K-means algorithm can be tailored to the specific application. A modification of the original technique to quantization within a windowing system is presented. The proposed approach is able to account for the system's shared colour map so display of multiple images is improved.

Bibliography

- [BA91] R. Balasubramanian and J. Allebach. A new approach to palette selection for color images. *Imaging Technology*, 17:284–290, December 1991.
- [Bru65] J. D. Bruce. Optimum quantization. Technical Report 429, Massachusetts Institute of Technology, 1965.
- [Bru77] P. Brucker. *On the complexity of clustering problems*, pages 45–54. Springer-Verlag, New York, NY, 1977.
- [CCW92] C. A. Chaudhuri, W. T. Chen, and J. S. Wang. A modified metric to compute distance. *Pattern Recognition*, 7(25):667–677, 1992.
- [CFM93] P. Coltelli, G. Faconti, and F. Marfori. On the application of quantization and dithering techniques to history of arts. *EUROGRAPHICS*, 12:351–362, 1993.
- [Cri92] John Cristy. quantize, 1992. ImageMagick Image Processing Software Package.
- [Dek94] Anthony H. Dekker. Kohonen neural networks for optimal colour quantization. *Network: Computation in Neural Systems*, 5:351–367, 1994.
- [FBF77] J. H. Friedman, J. L. Bentley, and R. A. Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software*, 3(3):209–226, September 1977.
- [FO89] Eugene Fiume and Marc Oulette. On distributed probabilistic algorithms for computer graphics. *Proceedings of Graphics Interface*, pages 211–217, 1989.
- [Fou94] Alain Fournier. From the look of things. *Proceedings of Graphics Interface*, pages 157–164, 1994.
- [Fri93a] B. Fritzke. A growing and splitting elastic network for vector quantization. In *Neural Networks for Signal Processing: Proceedings of the 1993 IEEE-SP Workshop*, pages 281–290, Maryland, 1993.
- [Fri93b] B. Fritzke. Vector quantization with a growing and splitting elastic net. In *ICANN'93: International Conference on Artificial Neural Networks*, pages 580–585, Amsterdam, 1993. Springer.
- [FS75] R. Floyd and L. Steinberg. An adaptive algorithm for spatial gray scale. *Society of Information Display, Digest of Technical Papers*, pages 36–37, 1975.

- [FvDFH90] James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics, principles and practice*. Addison-Wesley, 1990.
- [GJW82] M. R. Garey, D. S. Johnson, and H. S. Witsenhausen. The complexity of the generalized lloyd-max problem. *IEEE Transactions on Information Theory*, 28:255–256, March 1982.
- [GP88] M. Gervautz and W. Purgathofer. A simple method for color quantization: Octree quantization. In *New Trends in Computer Graphics*, pages 219–231. Springer-Verlag, New York, NY, 1988.
- [Gra84] R. M. Gray. Vector quantization. *IEEE ASSP*, 1:4–29, April 1984.
- [Hal89] Roy Hall. *Illumination and Color in Computer Generated Imagery*. Springer-Verlag, New York, NY, 1989.
- [Hec80] Paul Heckbert. *Color Image Quantization for Frame Buffer Display*. Bachelor's thesis, Architecture Machine Group, MIT, 1980.
- [Hec82] Paul Heckbert. Color image quantization for frame buffer display. *ACM, Computer Graphics*, pages 297–304, 1982.
- [HN90] Robert Hecht-Nielsen. *Neurocomputing*. Addison-Wesley, 1990.
- [Hod88] Michael H. Hodgson. Reducing the computation requirements of the minimum stance classifier. *Remote Sensing of Environment*, 25:117–128, 1988.
- [JX93] Gregory Joy and Zhigang Xiang. Center-cut for color-image quantization. *The Visual Computer*, (10):62–66, 1993.
- [KKL90] Jari A. Kangas, Teuvo A. Kohonen, and Jorma T. Laaksonen. Variants of self-organizing maps. *IEEE Transactions on Neural Networks*, 1(1):93–99, 1990.
- [Koh91] Teuvo A. Kohonen. Self-organizing maps: optimization approach. In *Artificial Neural Networks*, pages 981–990. Elsevier, New York, 1991.
- [LBG80] Y. Linde, A. Buzo, and R. M. Gray. An algorithm for vector quantizer design. *IEEE Transactions on Communication*, 28(4):84–95, January 1980.
- [Mat92] V. John Mathews. Multiplication free vector quantization using l_1 distortion measure and its variants. *IEEE Transactions on Image Processing*, 1(1):11–17, January 1992.
- [MC92] V. John Mathews and S. H. Cho. Improved convergence analysis of stochastic gradient adaptive filters using the sign algorithm. *IEEE Transactions on Acustics, Speech and Signal Processing*, 35:450–454, April 1992.
- [MG93] Jose L. Marroquin and Frederico Girosi. Some extensions of the k-means algorithm for image segmentation and pattern classification. Technical Report 1390, Massachusetts Institute of Technology, January 1993.

- [PBT93] Nikhil R. Pal, James C. Bezdek, and Eric C-K Tsao. Generalized clustering networks and kohonen's self-organizing scheme. *IEEE Transactions on Neural Networks*, 4(4):549-557, July 1993.
- [Pos91] Jef Poskanzer. ppmquant, 1991. PBM+ Image Processing Software Package.
- [SCB88] Maureen C. Stone, William B. Cowan, and John C. Beaty. Color gamut mapping and the printing of digital color images. *ACM Transactions on Graphics*, 7:249-292, October 1988.
- [SLP83] R. J. Stevens, A. F. Lehar, and F. H. Preston. Manipulation and presentation of multidimensional image data using the peano scan. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5(5):520-526, September 1983.
- [Tho90] Spencer W. Thomas. rlequant, 1990. Utah Raster Kit.
- [Uli87] Robert Ulichney. *Digital Halftoning*. The MIT Press, 1987.
- [VB95] Oleg A. Verevka and John W. Buchanan. Local k-means algorithm for color image quantization. *Proceedings of Graphics Interface*, pages 128-135, May 1995.
- [Ver95] Oleg A. Verevka. Color image quantization in windows systems with local k-means algorithm. *Proceedings of VI Western Computer Graphics Symposium*, pages 74-79, March 1995.
- [WPW90] S. J. Wan, P Prusinkewicz, and S. K. M. Wong. Variance-based color image quantization for frame buffer display. *Color Research and Application*, 15(1):52-58, February 1990.
- [Wu92a] Xiolin Wu. Color quantization by dynamic programming and principal analysis. *ACM Transactions on Graphics*, 11(4):348-372, October 1992.
- [Wu92b] Xiolin Wu. *Statistical Color Quantization for Minimum Distortion*. Springer-Verlag, New York, NY, 1992.
- [Yul67] J. A. Yule. *Principles of Color Reproduction*. John Wiley & Sons, Inc, 1967.