UNIVERSITY OF MALTA

MASTER OF SCIENCE IN ENGINEERING

Classification of pigments using Hyperspectral Imaging

Nathan MAGRO

Supervised by
Dr Alexandra BONNICI

Co-supervised by
Dr Stefania CRISTINA

A dissertation submitted in partial fulfilment of the requirements for the degree of Master of Science in Engineering

by the

Faculty of Engineering

October, 2021

Copyright Notice

Classification of pigments using Hyperspectral Imaging

Abstract

Hyperspectral imaging (HSI) is used in art conservation to capture signal differences in wavelength not visible to the human eye. This work investigates the use of HSI in the range of 400nm to 1000nm, to segment an image of a painting into regions consisting of different paint characteristics and identify the paints within the region.

Image segmentation algorithms often need to balance the trade-off between under-segmnetating the image, that is, grouping dissimilar regions, and over-segmentation, that is, fragmenting a region into separate regions. This dissertation proposes a Spectral Similarity Merging algorithm which prioritises homogeneity of regions, by keeping a low intra-class variance, while reducing over-segmentation. This was possible by modifying the Kernel Spectral Angle Mapper (KSAM) similarity metric into Kernel Spectral Correlation Mapper (KSCM) which enhanced the merging properties with regards to homogeneity. In fact, Spectral Similarity Merging (SSM) yielded an F-score of 48.5%, 39.8% higher than Simple Linear Iterative Clustering (SLIC), and an over-segmentation of 67.3%, 27.5% less than SLIC.

Following the segmentation problem is the pigment classification method is carried out to identify the paints that constitute a painted region. This work proposes three methods, namely (i) a Hierarchical Paint Analysis (ii) a Global Non-negative Matrix Factorisation (NNMF) Paint Analysis, both of which were developed to attempt paint mixture identification and which are limited to two paints per mixture, and (iii) a Direct Classification of Paint Mixtures, which was developed to compare the performance of KSCM and NNMF when the mixtures are present in the reference library. The results obtained show that the Global NNMF Paint Analysis performs better than the Hierarchical Paint Analysis, producing an average accuracy of 76%. For the third method, The Direct Classification of Paints by KSCM yielded an accuracy of 84%, 14% more accurate than NNMF, meaning that KSCM is better when reference spectra are available, with the additional advantage of being less susceptible to spectral variance.

Acknowledgements

I would like to thank my supervisors Dr. Alexandra Bonnici and Dr. Stefania Cristina for their consistent support and guidance for the duration of this dissertation. I would also like to thank the University of Malta, and Department of Systems and Control Engineering for the opportunity to participate in this study. Lastly, my deep I would like to thank my wife and family for their constant support and encouragement.

Contents

Copyright Notice i

Abstract ii

Acknowledgements iii

Contents iv

List of Figures vii

List of Tables x

List Acronyms xi

List Abbreviations xiii

1 Introduction 1
- 1.1 Research Focus . . . 2
- 1.2 Aims and Objectives . . . 3
- 1.3 Contributions . . . 3
- 1.4 Dissertation Organisation . . . 4

2 Background 5
- 2.1 Electromagnetic Spectrum . . . 5
  - 2.1.1 Reflectance, Transmittance and Absorbance . . . 6
- 2.2 Materials . . . 7
  - 2.2.1 Pigments . . . 9
  - 2.2.2 Binders . . . 10
  - 2.2.3 Surfaces . . . 10
  - 2.2.4 Pigment Mixtures . . . 11
- 2.3 Imaging Equipment . . . 13
  - 2.3.1 X-ray Fluorescence . . . 14
  - 2.3.2 Raman Spectroscopy . . . 14
  - 2.3.3 Fibre Optic Reflectance Spectroscopy . . . 15

2.3.4 Multispectral Imaging 16
2.3.5 Hyperspectral Imaging 16
2.4 Summary 19

3 Literature Review **20**
3.1 Imaging Techniques 20
3.1.1 Raman Spectroscopy 20
3.1.2 X-ray Fluorescence 21
3.1.3 Fibre Optic Reflectance Spectroscopy 21
3.1.4 Multispectral Imaging 22
3.1.5 Hyperspectral Imaging 22
3.2 Pigment Processing Techniques 23
3.2.1 Filtering 23
3.2.2 Feature Extraction 23
3.2.3 Pigment Identification 28
3.3 Summary 33

4 Methodology **35**
4.1 Materials and Equipment 35
4.2 Data Acquisition 38
4.3 Filtering 40
4.3.1 Spatial Domain Filtering 41
4.3.2 Filtering Pixel Vectors 44
4.4 Feature Extraction 45
4.4.1 Paint Segmentation 45
4.5 Paint Classification 53
4.5.1 NNMF for Paint Analysis 55
4.5.2 Normalisation 56
4.5.3 Reflectance to KS Coefficients 57
4.5.4 Method 1 - Hierarchical Paint Analysis 59
4.5.5 Method 2 - Global NNMF Paint Analysis 64
4.5.6 Method 3 - Direct Classification of Mixtures 66
4.5.7 Sensitivity to Variance of Spectra 68
4.6 Summary 68

5 Results and Discussion **70**
5.1 Segmentation Evaluation Methodology 70
5.1.1 Qualitative Evaluation Methodology 70
5.1.2 Quantitative Evaluation 73
5.1.3 Qualitative Evaluation Results and Discussion 75
5.1.4 Quantitative Evaluation Results and Discussion 82
5.2 Classification Evaluation Methodology 84
5.2.1 Hierarchical Paint Analysis - Evaluation Methodology 85
5.2.2 Global NNMF Paint Analysis - Evaluation Methodology 86
5.2.3 Direct Classification of Mixtures - Evaluation Methodology . . 87

5.2.4 Measuring Sensitivity to Variance of Spectra for KSCM and NNMF . . . 88
5.2.5 Ground Truth of Paints . . . 88
5.2.6 Hierarchical Paint Analysis - Evaluation Results and Discussion 88
5.2.7 Global NNMF Paint Analysis - Quantitative Evaluation Results and Discussion . . . 96
5.2.8 Direct Classification of Mixtures - Quantitative Evaluation Results and Discussion . . . 98
5.2.9 Sensitivity to Variance for KSCM and NNMF Results . . . 99
5.2.10 Comparing the Classification Methods . . . 100
5.3 Case Study . . . 103
5.3.1 SSM - Case Study . . . 105
5.3.2 KSCM vs Global NNMF - Case Study . . . 108

6 Conclusion **112**
6.1 Improvements and Further Work . . . 114

A Histograms for Pixel Offset **123**

B Classification results for other images **126**

List of Figures

2.1 The three phenomena that occur when a ray of light hits a surface. . . . 6
2.2 Characterisation of temperature magnitudes at different wavelengths. . . 7
2.3 The effects of different Pigment Volume Concentration (PVC) ratios. . . 8
2.4 The different surface reflections. 11
2.5 A block diagram showing the typical setup of X-ray Fluorescence Spectroscopy. 15
2.6 The different Hyperspectral Imaging (HSI) system setups 17

3.1 Comparing two different approaches to pigment analysis. 26
3.2 Illustration of how SAM calculates a similarity measure 30

4.1 Variance of a line scan against wavelength. 39
4.2 An illustration to show the effect of background subtraction on the data. 40
4.3 An illustration to show how quadrants of the Adaptive Kuwahara Filter may grow. 42
4.4 Plots to compare Median and Adaptive Kuwahara filters. 43
4.5 SNR at each wavelength of teh datacube for two ROIs. 44
4.6 A block diagram of the segmentation process. 47
4.7 Demonstration to show how superpixels having a high similarity should merge. 48
4.8 Showing the difference in intra-class variance of homogeneous and mixed paint regions. 49
4.9 Demonstration to show how superpixels that have high intra-class variance have their pixels re-assigned to neighbouring superpixels. 50
4.10 The process as seen by showing the results from each stage. 52
4.11 Illustrating the similarity between the two red pigments. 53
4.12 Two ROC curves for the two samples used to determine the right threshold. 54
4.13 Illustrating the learned KS coefficients for the nine paints in the dataset. 59
4.14 Simulation of a mixture of paints using the KM theory. 60
4.15 An illustration for the paint reflectance spectra present in the spectral library. 61
4.16 A block diagram that illustrates the pipeline for the classification process of a segmented image. 62

4.17 Pigments spectra separated by their class to show the difference in spectral variance for each class. The Cadmium Red Hue and Permanent Geranium Lake red paints have a very similar signature. On the other hand, although the Viridian Green and Emerald Green paints seem to be very similar, the difference in the location of their reflective peaks is sufficient to allow distinction between the two pigments. 63
4.18 Illustrating the features and weights obtained after performing NNMF for a rank of 2. 66
4.19 An example to show how the weights can be distributed by an NNMF of rank 3. 67

5.1 The results for Region Merging stage of the segmentation process. . . . 75
5.2 The segmentation result for the Pixel Re-assignment stage, sub-routine 1. 76
5.3 The segmentation result from the Pixel Re-Assignment, sub-routine 2. . 77
5.4 The segmentation result for the Checking for Further Merging stage. . . 78
5.5 Comparing different segmentation techniques for the task of paint segmentation. 78
5.6 Results for four different samples images. 80
5.7 A case where paint separation fails. 81
5.8 Varying the initial region merging threshold to check for improvements. 81
5.9 Pixel Distance Probability Density plot for all mock-up samples. 84
5.10 A box plot to compare over-segmentation and F1-score improvements. . 85
5.11 Paint RGB images. 89
5.12 Paint ground truths for layered paint images. 90
5.13 Paint ground truths for mixed paint images. 91
5.14 Classification of each feature in the form of binary labels. 92
5.15 Two examples of symmetrical matrices obtained from NNMF in the Hierarchical Paint Analysis method. 94
5.16 Demonstration to show that layered paints are harder to distinguish as mixtures when the top paint layer is highly reflective. 95
5.17 Classification result of Image 3 when separating the paint regions by paint class. 97
5.18 Classification result of Image 3 when separating the paint regions by individual paints. 98
5.19 Classification result of Image 3 for KSCM. 100
5.20 Classification result of Image 3 for NNMF. 101
5.21 Similarity between a mixture of two paints and the individual paints used in the mixtures. 103
5.22 RGB representation of ‘The Birth of John the Baptist’ and the ground truth signatures. 104
5.23 A snapshot example of a fine detailed region. 105
5.24 Segmentation results on case study. 106
5.25 Results obtained for a fine detail region. 106
5.26 A snapshot example of a region painted with similar paints. 107
5.27 Results obtained for a red region containing multiple red paints. 107

5.28 Classification results obtained from case study. 108
5.29 KSCM classification results for blue and red paints. 109
5.30 Illustrating the difference between the red paint reference number 5 and actual signature in misclassified region. 110
5.31 NNMF classification results for blue and red paints. 111

List of Tables

3.1 Summary of literature focusing on the type of equipment and methods used. 33

4.1 The mass to mass paint mixing ratios for the homogeneously paint mixture samples with Paint 1 (50%) and Paint 2 (50%). 36
4.2 The mass to mass paint mixing ratios for the homogeneously paint mixture samples with Paint 1 (60%) and Paint 2 (40%). 37
4.3 The mass to mass paint mixing ratios for the homogeneously paint mixture samples with Paint 1 (40%) and Paint 2 (60%). 37
4.4 Noting the different elements used to create the different paints 38

5.1 Over-Segmentation (OS) and Under-Segmentation (US) results for SSM when compared to K-means and SLIC. 83
5.2 F1-Score results for SSM when compared to K-means and SLIC. 83
5.3 Comparing accuracy results for KSCM, KSAM, SCM and SAM for the problem of separating mixed regions from individual regions. 93
5.4 Comparing accuracy results for KSCM, KSAM, SCM and SAM for the problem of correctly identifying the pure paints in individual regions. . 93
5.5 NNMF metric results for various homogeneous paint mixtures, separating them by class. 94
5.6 NNMF metric results for various layered paint mixtures, separating them by class 95
5.7 Accuracy results for Global NNMF Paint Analysis method on both classification by class and by individual paint. 96
5.8 Accuracy results for Direct Classification of Paint Mixtures. 99
5.9 Sensitivity to variance for KSCM and NNMF results for both KSCM and NNMF, when applied to a 12-Paint reference dataset that only includes mixed paints, and a 22-Paint reference dataset that also includes the individual paints. 102

List of Acronyms

KSAM Kernel Spectral Angle Mapper. ii, 28, 31, 60, 62, 68, 86, 89–91, 95, 113

KSCM Kernel Spectral Correlation Mapper. ii, vi, x, 66–68, 71, 75, 77, 85–89, 94, 95, 98–100, 102, 103, 108, 110, 113, 114

SSM Spectral Similarity Merging. ii, 45, 53, 70, 72, 73, 78, 79, 82, 84–86, 89, 105–107

SLIC Simple Linear Iterative Clustering. ii, 25–27, 45, 50, 72, 79, 82–84, 105–107, 113, 114

NNMF Non-negative Matrix Factorisation. ii, v, vi, viii, x, 29, 32–35, 55, 60, 63–68, 85–88, 93, 98–100, 102, 103, 108, 110, 114

PVC Pigment Volume Concentration. vii, 8, 9, 21

HSI Hyperspectral Imaging. vii, 2, 3, 14, 16–19, 22–26, 32, 33, 39, 45, 112

XRF X-Ray Fluorescence. 2, 14, 15, 20–22, 33

RS Raman Spectroscopy. 2, 14, 15, 20–22

FORS Fibre Optic Reflectance Spectroscopy. 2, 14–16, 21–23

VIS Visible. 2, 6

SAM Spectral Angle Mapper. 3, 28–31, 34, 47, 62, 86, 90, 91, 95, 113

SCM Spectral Correlation Mapper. 3, 28, 30, 47, 50, 62, 68, 86, 95, 113

NIR Near Infrared. 6, 9, 19

KM Kubelka-Munk. 12, 29, 33, 35

MSI Multispectral Imaging. 16, 21, 22

DBSCAN Density Based Spatial Clustering of Applications with Noise. 24, 25

WT Watershed Technique. 24

ED Euclidean Distance. 27

ICA Independent Component Analysis. 29, 32

SISAL Simple Identification Split Augmented Lagragian. 31

MVES Minimum Volume Enclosing Simplex. 31

MVSA Minimum Volume Simplex Analysis. 31

NMF-MVT Non-negative Matrix Factorisation - Minimum Volume Transform. 31

PCA Principal Component Analysis. 32

SNR Signal to Noise Ratio. 38, 41

VNIR Visible-Near Infrared. 61

List of Abbreviations

LY Lemon Yellow. 94, 95, 102

CRH Cadmium Red Hue. 94, 95, 102

PR Permanent Rose. 94, 102

VG Viridian Green. 94, 95, 102

CBH Cerulean Blue Hue. 94, 95, 102

PGL Permanent Geranium Lake. 94, 95

EG Emerald Green. 94, 95, 102

PB Phthalo Blue. 94, 95, 102

FU French Ultramarine. 94, 95, 102

Chapter 1

Introduction

Historical works of art, whether in the form of paintings, sculptures or pottery fragments, are of value not just for artistic content, but also for their cultural value. For example, works of art, such as those from the different Chinese Dynasties depict not only long-lost scenes but also serve to give insight to past cultures [1]. Therefore, the preservation of these artefacts is considered essential for the protection of cultural heritage. Unfortunately, these works of art can suffer from degradation due to careless human handling and environmental factors and thus, restorative studies play a necessary role in the conservation process.

The focus of this dissertation will be paintings. Here, organic materials, such as linseed oil, are mixed with the powdered form of the pigment to create paint. The powdered form of a pigment is typically made of the right consistency of minerals to create the specific paint. In the past, most of the paints consisted of a single mineral, however, more recent paints can have up to three different minerals. For example *raw amber* consists of the minerals iron oxide and manganese oxide. The organic materials that are mixed in with the pigments may further effect the appearance of the painted region. For example, in the case of oil-based paints, an effect called yellowing may start to develop over time and thus, art restorers need measures to help identify the true colours and properties of different paint samples.

Analysis of paintings can be carried out either through invasive or through non-invasive techniques. Invasive methods typically include the destructive sampling of the paint for chemical analysis in laboratories. Non-invasive methods, on the other hand, make use of imaging and other scanning methods which are non-damaging to the painting. Thus, non-invasive methods are generally preferred over their invasive counterparts, whereas invasive methods can be used when a small location of the painting requires restoration, such as known locations of small cracks [1–3].

Non-invasive procedures include various forms of spectroscopy that range from X-Ray Fluorescence (XRF) to Hyperspectral Imaging (HSI). Methods such as Raman Spectroscopy (RS) and Fibre Optic Reflectance Spectroscopy (FORS) provide spectral information for a single point, whereas image-based methods such as XRF and HSI provide spectral-spatial information at the cost of having a smaller range of wavelengths. This spatial information can be very important to a restoration process because it can provide a guideline for borders of paints or of regions painted with the same paint. Such information is not captured in point-based spectroscopy. Moreover, comparing XRF to HSI, less radiation is being imposed on the painting by using HSI as the spectrometer is only capturing light reflected off the surface instead of radiating X-rays on the surface of the painting. What is more interesting about image-based methods is that, analysis can result in discoveries of features that are hidden to the human eye [3].

1.1 Research Focus

The most important part of a restoration process is the correct identification of materials on a surface. The materials are usually a mixture of two components: a pigment and a binder. The pigment is the component that provides the colour, whereas the binder gives the desired consistency to the solution. Organic materials such as linseed oil do not interfere with the identification of paints in the Visible (VIS) range because they lack absorption bands in this range. When a paint is mixed with another paint in a homogeneous manner to create a different shade, the identification process becomes harder, because an artist would need to know which paints were used to create the mixture.

In this study, the HSI method is used to analyse mock-up samples, exploiting the spatial-spectral property of such spectroscopy method to address a superpixel segmentation algorithm, which aims to split the image into large segments by colour or spectral difference, such that areas with different paints can be separated. Existing segmentation techniques mentioned in Chapter 3 can already produce such results, however, at the expense of fragmenting regions into multiple segments even if they have the same spectral features. This is also known as over-segmentation, and the aim is to reduce this such that regions having homogeneous spectral properties will not be fragmented. We will also be focusing on labelling these regions according to what paint they contain, and identifying paints in a mixture. A comparison between different classification techniques will be done.

1.2 Aims and Objectives

In art restoration, it is important to be as accurate as possible to conserve the historical value of a painting. One must keep in mind that each time a painting undergoes a restoration process, the painting will be at risk due to handling. Therefore, analytical methods must be accurate and precise such that processes will not need to be repeated. In this study, the aim is to develop paint analysis processes based on HSI focusing on oil-based paints, and adhere to the following objectives:

- To create mock-up samples and acquire the data by using a HSI device.
- To investigate a superpixel segmentation algorithm, with the aim to reduce over-segmentation without compromising the boundary adherence property.
- To identify paints within mixtures using a method of spectral unmixing.
- To develop, modify or extend a classification algorithm that classifies the pure materials at a high accuracy rate.
- Robust and perform well on all mock-up samples created.
- To carry out a case study to test the process in real world scenarios.
- To gather knowledge in pre-processing algorithms, signal processing and classification algorithms, and additionally, gather general knowledge regarding the restoration of artworks using non-invasive methods.

1.3 Contributions

This study contributes towards the art restoration community by providing i) a process that is capable of separating regions in a painting by paint homogeneity, ii) a paint identification method, and iii) an analysis of paint mixtures. For the former contribution, a study has been found where the researchers were capable of separating regions of a painting by paint features. However, their method is restricted by the regular sizes of these regions, and although the homogeneity property still stands, some of the regions can be merged together [4]. Therefore, this study shall reduce the over-segmentation without compromising boundary adherence. On the other hand, the second contribution aims to provide a paint identification method that is better than existing metrics, such as Spectral Angle Mapper (SAM) and Spectral Correlation Mapper (SCM) [5]. The final contribution provides a process that helps in understanding the mixing of two paints.

Here, an identification of two paints is computationally obtained and compared with the ground truth. This contribution should aid restorers in creating the right mixtures of paints with higher accuracy.

1.4 Dissertation Organisation

The rest of this dissertation is organised as follows; Chapter 2 presents a background on the materials and equipment; Chapter 3 presents a thorough literature review on recent developments in this field; Chapter 4 states the chosen materials and data acquisition methods; Chapter 5 includes the methods and algorithms applied; Chapter 6 discusses the results obtained; and Chapter 7 concludes on this work.

Chapter 2

Background

In this chapter, the basics of the electromagnetic spectrum, especially that in the visible and near infra-red regions will be discussed. Moreover, the basic composition of pigments, binders and their mixtures will be discussed such that a good understanding is obtained. Finally, some background on the equipment that is used in this field is given such that the main differences, advantages and limitations are understood.

2.1 Electromagnetic Spectrum

All radiation in the electromagnetic spectrum can be represented by a single term, photon. The energies that these photons carry can be absorbed by atoms representing matter. When these atoms absorb some of these photons, they go through different phenomena depending on the energy striking them. Due to the conservation of energy, the energy being absorbed is converted to other energies such as light and heat. Lifeforms are capable of seeing light because their light sensors are sensitive to these energies that atoms release. Human eyes, for instance, are sensitive to three wavelengths in the visible spectrum; red, green and blue [6]. In the visible region, electrons absorb photons and go through a process called transition. In this process, the atoms move to a higher state of energy which is dependent of the level of energy they are capable of absorbing. When these atoms absorb this level of energy they become unstable, forcing them to release energy in order to return to their original stable state. The released energy is what causes humans to see different colours on different surfaces. In the infrared region, these atoms also experience changes, however, these energies are not enough to cause the atoms to transit to a higher state. Instead, these atoms experience vibrations within their bonding structure [7, 8]. This would lead us to better understand the three main phenomena: reflectance, transmittance and absorbance of when light interacts with matter.

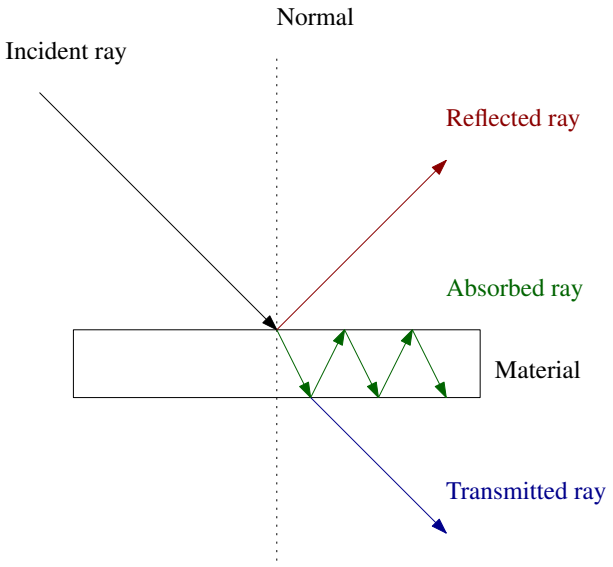

FIGURE 2.1: The three phenomena that occur when a ray of light hits a surface. The ray of light can be reflected, absorbed or transmitted.

2.1.1 Reflectance, Transmittance and Absorbance

Reflectance is the amount of light reflected back from a surface, but depending on the surface, incident light intensity may not be equal to reflected light intensity. On the other hand, transmittance is the amount of light that passes through the surface and exits the material from the other end. Finally, absorbance is the amount of light that is absorbed by the material. Figure 2.1 shows how the three phenomena could occur when an incident ray of light reacts with a material. There are two ways that can cause light reflectance intensity to change when light is incident on a surface, either by transmittance, absorbance, or the mixture of both, and the relation between the three is given by Equation 2.1.

$$R = 1 - (T + A) \tag{2.1}$$

where T is transmittance, R is reflectance and A is absorbance. Although the combination of the three is what gives the colour that we see with our naked eyes, for opaque surfaces transmittance is usually ignored. Additionally, absorbance is dependent on the thickness of the surface layer and particle sizes of the material.

One may also argue that because there is absorbance, there is also the presence of emittance. However, emittance is energy radiated by the surface after the energy from the incident ray is absorbed by the surface. This type of energy is released at a much longer wavelength than that of the Visible (VIS) or Near Infrared (NIR), between 8-14 μm or long-wave IR. Francisca et al. [9] showed that surfaces at temperatures between 200°C

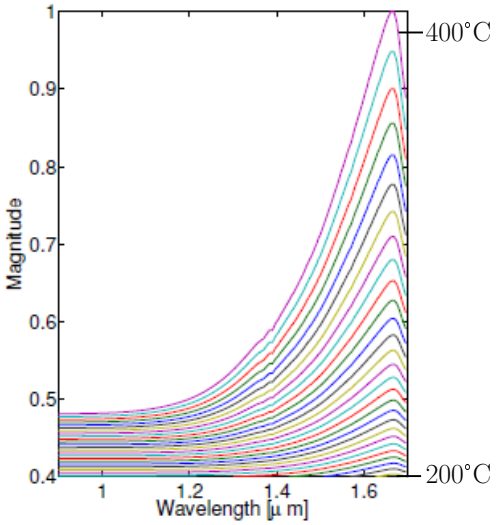

FIGURE 2.2: Characterisation of temperature magnitudes at different wavelengths for multiple temperature values ranging from 200°C - 400°C at 5°C intervals reproduced from [9]

and 400°C result in different intensities. In fact, at the 200°C mark, the intensity is almost flat as shown in Figure 2.2. This means that for surfaces at a very low temperature, such as room temperature, the effect of emmitance can be assumed to be negligible if the environment is in thermal equilibrium.

2.2 Materials

The type of restoration a painting requires is highly dependent on the materials used by the artist. Materials could be surfaces, pigments and the binding solutions. In the past, most of the powdered pigments used were created by crushing minerals, such as azurite and malachite, among others, whereas, most binding solutions are organic substances made from animals or plants. When both materials are mixed together a viscous solution is created, providing the artist with different properties. Such properties include the difference in consistency for different effects, resulting in various reflective properties. The reason why a binder changes the reflective properties of a pigment is because the abundance of powder in the solution reduces, causing light energy to be reflected back before it hits powder particles, as shown in Figure 2.3. Properties that may affect the reflectivity of a solution include: refractive index; particle size; and absorption coefficients.

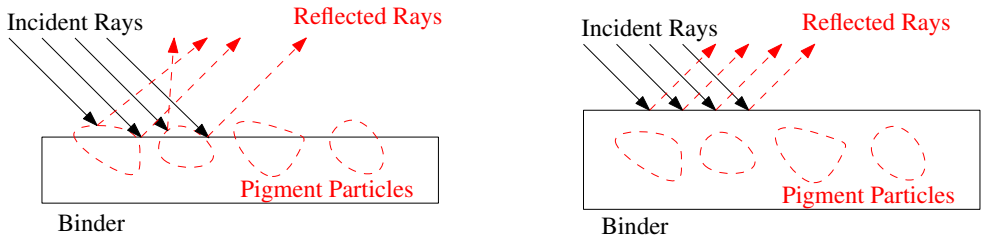

(a) A good PVC Ratio (b) A low PVC Ratio

FIGURE 2.3: Effects of different Pigment Volume Concentration (PVC) ratios which show how the reflected rays are less affected by the binding solution in (a).

The refractive index of a material, calculated by Equation 2.2, is a number that describes the change in the speed of light in a substance or material:

$$n = \frac{c}{v} \tag{2.2}$$

where c is the speed of light in vaccuum and v is the phase velocity of light in the material [10]. From the refractive index, a reflectance value can be measured if the refractive index of both media is known. This is given by Equation 2.3:

$$R_0 = \left| \frac{n_1 - n_2}{n_1 + n_2} \right|^2 \tag{2.3}$$

where n_1 is the refractive index of the first medium and n_2 is the refractive index of the second medium [10]. The two mediums can be either air and paint, or a binder and paint.

In Equation 2.4 the formula for the absorption coefficient is given. It measures the loss of radiant flux in a medium. When added with a scattering coefficient, a total attenuation coefficient is measured, which measures the loss of radiant flux due to scattering:

$$\mu = \mu_a + \mu_s \tag{2.4}$$

where μ is the total attenuation coefficient, μ_a is the absorption coefficient and μ_s is the scattering coefficient.

Particle size directly relates to the size of the molecules in the pigments. It can be measured either by volume or by surface area [11]. As Gueli et al. mentions, the particle

size can affect the reflectivity of a material by a considerable amount.

PVC is a property that describes the pigment to binder ratio:

$$PVC = \frac{\text{PigmentVolume}}{\text{PigmentVolume} + \text{BinderVolume}} \tag{2.5}$$

In [12], Sands argues that the best mixture is when the binder fills all the air pockets between the particles of the pigment, the Critical-PVC, at which the characteristics of the pigment are not lost while characteristics of the binder are gained. Sands also shows how an overdose of the binder - lower than the CPVC - will result in loss of characteristics of the pigment because the pigment particles will be much less dense. This means that light rays will mostly reflect off the binder and not the pigment particles. Figure 2.3 demonstrates how this occurs. It is important to note that for the visible spectrum, binders such as linseed oil do not minimally reflect visible light - they are transparent - and therefore, one should still be able to see the pigments through the overflowing binder.

A more detailed review of different pigments and binders is given in the following subsections, specifically the points of interest on the spectrum of the pigments and the properties of binders.

2.2.1 Pigments

Humans see colours because each surface is able to absorb and reflect different wavelengths, hence, creating a colourful environment. However, this can only be possible by the prior mentioned properties that each surface holds, which are effects of the very different minerals present in the media. Pure pigments are usually dominated by a specific mineral. To create different shades of pigments, artists used to combine different minerals together.

Blue pigments are usually made up of azurite or cobalt minerals [13], which are both copper based. Although different blue pigments have different absorption coefficients and absorption bands, most blues have a reflective peak at about 450nm followed by an absorption band which rises slowly towards the NIR region. Greens are also mostly made up of copper elements. Malachite and Verdigris are both examples of copper based green pigments [13]. These have a reflective peak around the 550nm wavelength, which is also followed by an absorption band that slowly rises towards the NIR region. On the other hand, reds start with an absorption band that rapidly rises to peak at around 650nm, which stays at a high reflective value at the NIR region as well. The oldest red pigments are found to be made up of the calcium based mineral, hematite. Other red pigments can be made up of oxide based minerals, such as ochres [1].

Although the range of colours follow the general rule of their respective absorption and reflective ranges, due to the different mixture of minerals present in each range of colours, different reflective spectral curves are created which is why signal processing techniques can be used to classify pigments in a supervised manner.

2.2.2 Binders

Pigments are usually presented as a powdered medium. In order to transform this powder into a medium that can be spread on a surface, for example canvas, artists need to mix the powder with a binding agent. There are various kinds of binding agents, namely: oils, proteins and carbohydrates; among others [14]. These binders present different properties, not only to the naked eye, but also different properties in deterioration and in reflectivity. Ropret et al. studied the effects of ageing on binders. It was found out that carbohydrates, such as gum arabic, start to peel off when aged. Ropret et al. also argue that oils were found to lose colouration due to ageing, a process typically referred to as yellowing. On the other hand, protein based binders such as egg yolk, break down when exposed to UV light. Additionally, moisture causes a slow hydrolysis process between protein links. The paper also gives an interesting view to the spectral profiles of the samples when aged by different methods.

2.2.3 Surfaces

One of the main factors that affect light behaviour is the smoothness of a surface. There are three kinds of surfaces that affect light rays differently; shiny surfaces, transparent surfaces and dull surfaces. Dull surfaces are those whose surface smoothness factor is small. They allow light rays to undergo the process of diffuse reflection, where, a light ray is dispersed equally across all directions above the surface. These surfaces allow equal identification of brightness levels and colour from any perspective. Shiny surfaces on the other hand, reflect light rays at the same angle of the incident ray to the normal [16]. Both of these phenomena are illustrated in Figure 2.4. Surfaces that have a transparency index are more complex. These surfaces undergo reflection, refraction, and diffuse reflection from the surface behind the transparent surface. For example, a glass window is hard to see while the scenery behind the window is visible in good definition, but, a translucent window is more visible while the scenery behind it would appear to be blurred out [16].

In this work, the type of surfaces used are stretched canvases. They are dull white surfaces typically painted in a titanium white primer that allow pigment solutions to dry

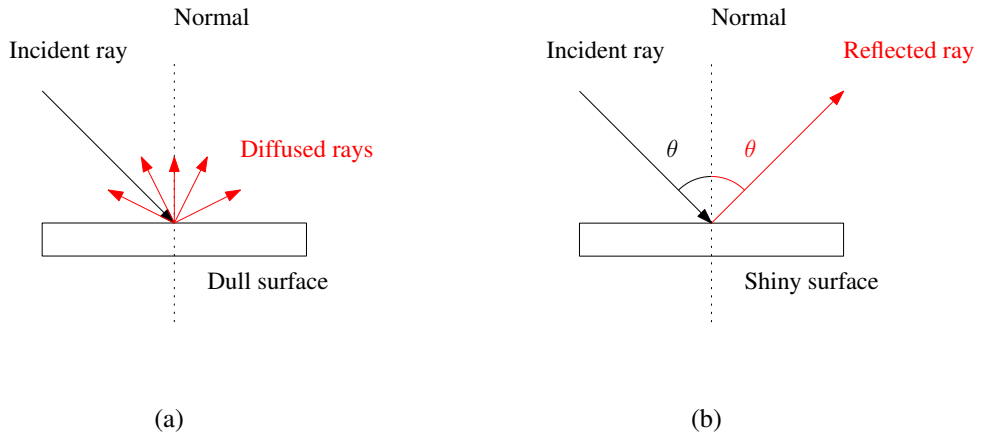

FIGURE 2.4: Surface reflections; (a) Diffuse reflection on a dull surface, (b) Specular reflection on a shiny surface

in place. The reason why they are painted in white is to allow for maximum reflectivity of the paint layer. Therefore, since we will be working with diffuse reflections, data acquisition will be totally dependent on the proper lighting of the samples.

2.2.4 Pigment Mixtures

Pigments are typically mixed with a binding solution such that the artist could create paint of specific properties. When a pigment is mixed with a binder, the different particles become homogeneously mixed. This creates a difference in the reflectivity of the pigment as stated earlier in the beginning of this section. This difference happens because a scattered ray of light leaving a pigment particle has the chance to interact with a different pigment particle which has different absorbance and scattering properties [17]. Moreover, artists also mix pigments together to create various shades at their own will. This results in much more complex reflectance curves as the signatures will represent various pigments, and follow a subtractive model, if studied as reflectance. There are two main types of pigment mixtures. The first type of mixture is formed when two pigments are homogeneously mixed together. This type of mixture would typically create another shade of pigment dependent on the two pigments used. For example, a blue and a yellow would create a green shade. Of course, if a restorer is looking at this green shade by the naked eye they would not know that it was originally made by mixing two other pigments. Therefore, the importance of pigment analysis in such situations.

The second type of pigment mixtures is produced when two pigments are layered on top of each other. These types of mixtures would also create different shades. However, these may be harder to interpret by the naked eye, since the patch will appear to be darker because the pigment in the under-layer is absorbing light as well. Moreover, in some cases the pigment on the top-layer will have enough thickness to give the analyst an idea of what pigment it may be. The layered form of mixtures are easier to analyse using a destructive approach because two layers can be separated under a microscope. However, it would be more ideal if layered mixtures can also be determined using computational models. Such models include complex mathematical solutions, and prior knowledge on the layer thicknesses would be required to have a non-invasive study, such as done in the Kubelka-Munk Theory [17]. On the other hand, if the pigments in a homogeneous solution are known, a ratio of the pigments in the mixture can be measured [18].

2.2.4.1 Kubelka-Munk Theory

The Kubelka-Munk (KM) theory was developed to measure the absorption a and scattering s coefficients of a mixture of pigments by knowing their concentrations c, or to measure these coefficients when the mixtures are known. The theory can be used for turbid (opaque) materials. By having a library of these coefficients, a signature can be produced which can be compared to the unknown mixture. There are two versions of this theory for homogeneous mixtures, the single-constant and the two-constant KM theories. Starting by the two-constant model, both pigments in the concentrated mixture have two independent variables, that is the absorption and scattering coefficients are independent, given by Equation 2.6:

$$\left(\frac{K}{S}\right)_{\lambda,m} = \frac{c_1 a_{\lambda,1} + c_2 a_{\lambda,2} + ... + c_n a_{\lambda,n}}{c_1 s_{\lambda,1} + c_2 s_{\lambda,2} + ... + c_n s_{\lambda,n}} \tag{2.6}$$

where K/S is the ratio of absorption and scattering, a is the absorption coefficient, s is the scatter coefficient, c is the concentration value, n is the number of pigments in the mixture and λ is the wavelength [19]. On the other hand, in cases where the scattering coefficient is very strong, such as when mixed with white pigments, a single-constant model can be applied. In the single-constant model, the scattering coefficient can be assumed to have a constant value of 1 throughout all wavelengths, but with the inclusion of the tint (the white pigment). Therefore, the model for the single-constant is given by Equation 2.7:

$$\left(\frac{K}{S}\right)_{\lambda,m} = \frac{c_1 a_{\lambda,1} + c_2 a_{\lambda,2} + ... + c_w a_{\lambda,w}}{c_w s_{\lambda,w}} \tag{2.7}$$

where w relates to the white pigment [17].

In cases where the coefficients of the pigments in a homogeneous layer are known, a different model can be applied to estimate the reflectance. The substrate ξ and layer thickness t are also required. This model is highly non-linear in nature and in certain cases, assumptions can be taken to simplify the computations [19]. The model is expressed by Equation 2.8:

$$R = \frac{1 - \xi \left(x - y \cdot \coth\left(yst\right)\right)}{x - \xi + y \cdot \coth\left(yst\right)} \tag{2.8}$$

where,

$$x = 1 + \frac{a}{s}$$
$$y = \sqrt{x^2 - 1}$$

and R is the total reflectance. For layers that do not scatter light, such as plastic films, one can assume that the scattering coefficient is 0, and for layers that fully scatter light, one can assume that the scattering coefficient is 1 [19]. This assumption is not usually preferred for paint because all paints scatter some light, and unless analysing pure Lead or pure Titanium white, the scattering coefficient will not be equal to 1 throughout all wavelengths. Similarly, one can assume that the media is turbid, hence the thickness of the material is infinite. This simplification does not hold for very transparent paint such as watercolours, but according to [20] it is valid for oil or acrylic paints.

One of the biggest limitations of such data is that, the wavelengths of the sample at hand must match the wavelengths that the scattering and absorption coefficients were measured with. If the wavelengths do not match, the simulated signatures will result in shifts, widening or narrowing of absorption or reflectance bands, which may result in inaccurate deductions of the pigment being analysed.

2.3 Imaging Equipment

All imaging techniques come from the basic idea of measuring light intensity values and converting them into specific digital values, which for us is depicted by a specific colour. The most basic form of conversion is the grayscale conversion, where each pixel can contain a value between 0 and 255. This is followed by an RGB camera which captures visible light (380-700nm) from three wavelength positions. These wavelengths are blue (448nm), green (550nm) and red (670nm), which are converted into a three valued vector [21]. Therefore, we can say that each pixel is a feature vector containing

three values. Although this vector may give information with regards to the pigment being analysed, it is often not enough to distinguish between pigments of a similar hue. Additionally, two identical colours under different lighting will appear to be different to the human eye. Hence, the importance of new imaging techniques that are capable of capturing the very details of these pigments such that proper distinction between very close pigments is made possible without the need of relying on invasive chemical methods. One must keep in mind that paintings can be damaged, up to a molecular level, if exposed to an extended duration of UV light through photochemical reactions [22]. Hence one requirement of these imaging techniques is to be fast when used, while at the same time a good resolution is also required, such that the analysis can be done at very fine detail. In the following subsections several different imaging spectroscopy methods, such as X-Ray Fluorescence (XRF), Raman Spectroscopy (RS), Fibre Optic Reflectance Spectroscopy (FORS) and Hyperspectral Imaging (HSI) are analysed.

2.3.1 X-ray Fluorescence

As the name XRF suggests, it works by emitting powerful x-ray particles onto the surface, as shown in Figure 2.5. When these waves hit the surface, they split up into fluorescence, Rayleigh scatter and Compton scatter, and some of the energy is lost through transmission [23]. Rayleigh scattering occurs when the wavelengths of the incident and scattered light are the same. On the other hand, Compton scattering happens when a photon is scattered with a gain in wavelength. Fluorescence happens when an energy beam hits the surface with the perfect amount of energy levels. If the energy levels are too high, the beam will just pass through the surface, whereas if the energy levels are too low, the beam will be absorbed with minimal production of fluorescence. However, if the energy is just right, that is the energy of the photon is slightly higher than the energy of the binding electron, then these electrons will be expelled from their shell, producing fluorescence [23]. This is how XRF can be used to determine pigments, by knowing which elements are present in the sample. Once, the elements are known, a library will help the analysts identify the pigment. One disadvantage of XRF that may be crucial in this field is that it cannot be used on oxides, and many pigments are made from oxides such as Viridian, a green pigment.

2.3.2 Raman Spectroscopy

The theory behind RS is very similar to that of XRF. It measures a Raman scattering of waves after a molecule is hit with a strong beam of light. While XRF measures

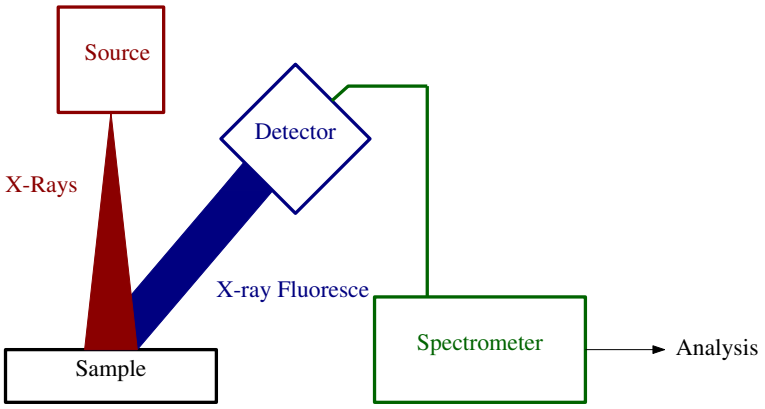

FIGURE 2.5: A block diagram showing the typical setup of X-ray Fluorescence Spectroscopy.

fluorescence off atoms, RS measures a raman scatter off molecules. Raman scatter is an inelastic scatter that does not only change the wavelength but also the phase. Due to the different light intensities of each method, they need to be used differently. Moreover, Raman scatter is much weaker than other phenomena that occurs when particles are struck by beams of light. In fact, impurities that may have a strong influence on fluorescence may affect Raman scatter negatively causing it to be harder to collect useful information [24]. Finally, Raman scatter cannot be used on metals or alloys and, as already seen in Section 2.2, most pigments contain metallic properties.

2.3.3 Fibre Optic Reflectance Spectroscopy

Another form of acquiring spectral data of a material is by using FORS, which returns a different kind of measure than RS or XRF because it does not use high energy beams of light such as X-rays or gamma rays but uses focused light techniques, such as a halogen lamp [24]. This is only possible by having a probe-like device capable of emitting a point-based light and capture the reflected signal. The reflectance depends on a set of variables as a function of wavelength: the refractive index of the material; the absorption coefficient; and the particle size of the atoms [3].

FORS works by shining light upon a surface whose reflected signal is passed through the fibre optic of the device. Hence, a reflective value of measurement will be obtained. The light passes through a filter within the device to create the desired wavelength in a sequence, such that, a complete signature with high resolution can be created, which represents the sample. One limitation of FORS is that it can only measure a reflectance value of a very small patch at a time, because of its physical limitation being a probe like device [24]. However, this can allow techniques to be very fast and versatile to read measurements across multiple points of interest. Also, due to the physical nature of the

sensors, the device typically returns very noisy data at the initial and final bands of its spectral range [24]. The FORS technique is typically used when points of interest are known and sampling of the entire painting is not required.

2.3.4 Multispectral Imaging

A more simple form of imaging is, in fact, Multispectral Imaging (MSI). The idea is to increase the number of filters inside the camera from three to ten or more [25]. An example of additional filter would be in the between the range of 700-1000nm, to capture infra red information of the surface. Of course, an additional filter will not provide enough information for pigment analysis, however, by having ten wavelengths of information, we can start to produce a geometric profile of the reflected light coming from the surface. In this particular case, when compared to the digital camera, one would have ten features rather than three, which is more than double the information.

2.3.5 Hyperspectral Imaging

Although MSI may provide enough information for pigment identification, it does not capture the finer details. Hence, complimentary methods are required to confirm results obtained from MSI, such an example is [26]. This requirement is less needed for HSI because the resolution is much higher and therefore specific features for each pigment or binder can be better identified. One could say that HSI alleviates the limitation of MSI providing richer data and information, at the cost of complexity and expense.

However, increasing the resolution automatically increases the required computational resources and computational time. Each pigment absorbs light at a different wavelength, which makes it impossible to create enough filters to satisfy each wavelength. It is also not very efficient to replace camera filters every time. Instead, HSI scanners use dispersive elements, such as prisms, to disperse the incoming ray of light into multiple wavelength specific rays [25]. Hence, by using high resolution HSI, we are able to image an entire region and capture all the useful information without any prior knowledge of the pigments on the surface.

2.3.5.1 Types of HSI Scanners

When it comes to historical artworks, specifically those on walls, the main problem is accessibility and portability. Therefore, three main types of HSI scanners have been developed; Whiskbroom, Pushbroom and Spectral, as illustrated in Figure 2.6. Each of

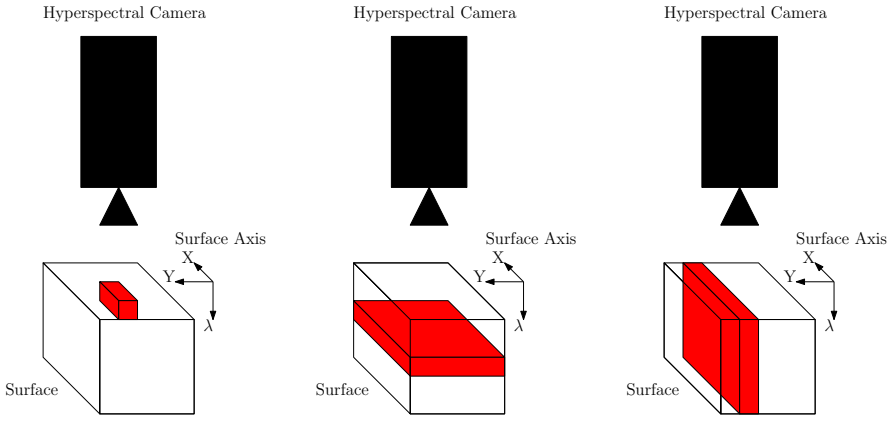

(a) Whiskbroom (b) Spectral (c) Pushbroom

FIGURE 2.6: The different HSI systems where the cube represents the dimensions of an image and the red part represents the points being scanned at one instance; (a) a Whiskbroom scanner, (b) a Spectral scanner, (c) a Pushbroom scanner. The red cross-section shows how each scanner scans the surface.

which operate in a different manner, but the end result is equivalent for all three methods [25].

The whiskbroom scanning method measures reflectance one pixel at a time, requiring only one detector element for each spectral band or wavelength. To do this, the detector measures a reflectance value, one wavelength at a time, of a light ray coming off the surface of a sample by using a dispersing element. Both the sample and the scanner must be stationary for this method such that only the detector element within the scanner moves. The whiskbroom design, therefore, is one of the slowest scanning methods and is used for high resolution imaging but requires perfect re-positioning for good repeatability [25].

The second method of scanning is the pushbroom design. Here, instead of measuring reflected light rays one pixel at a time, the detector measures a whole line of pixels at a time. This is possible by having a two-dimensional dispersing element. This design measures the entire spectrum of the light ray at once. It is mostly used in industrial applications where the scanner is in a fixed position and the sample is moving along an axis of the scanner. The performance of a pushbroom design is therefore much faster than the whiskbroom [25].

Finally, spectral scanning, is one of the simpler designs were a camera takes snapshots of a sample surface one wavelength at at time. This is achieved by installing band pass filters before the camera that alternate to create the spectral data cube. Area scanning is

one of the slowest scanning mechanisms and that is considered as a huge disadvantage. Therefore, it is used only when the required wavelengths are known, hence, adjusting or disabling any redundant filters [25].

2.3.5.2 Calibration

Just like any other sensing device, calibration is periodically required for HSI scanners. There are various factors that may affect measuring reflectance of a surface, one simple example is ambient light correction. There are three main calibration techniques that are required for every hyperspectral imager: wavelength calibration; radiometric calibration; and noise reduction or removal [27].

At manufacturing, calibration is an important step. Each type of scanner has their own calibration methods. Wavelength calibration is mostly applied to each scanner. In short, wavelength calibration assures that each array element of a detector points to the right wavelength, such that, the proper wavelength is assigned to each row of the data cube when analysing a sample [27].

In order to ensure that the readings are being interpreted well, a process called radiometric calibration is required. Here, a pixel-by-pixel correction is done by acquiring three images: reference image; sample image; and dark current image. The reference image is captured off a white reflective surface, usually of very high reflectivity, such as Spectralon [27]. As the name suggests, the sample image is the painting under investigation, whereas, the dark current image is an image of a completely covered lens. Hence, the percentage reflectance, used by most papers [26, 28, 29], can be calculated by Equation 2.9:

$$\boldsymbol{R} = \frac{S_\lambda - D_\lambda}{R_\lambda - D_\lambda} \times 100\% \tag{2.9}$$

where $\boldsymbol{R}$ is the percentage reflectance image, S_λ is the sample image, R_λ is the reference image, D_λ is the dark current image and λ is index for wavelength [27]. There also exists calibration for transmittance, however, it is not within the scope of this study as we only deal with reflectance. On the other hand, data normalisation is a far more important step, as it can also help correct ambient light effects. There are various ways of normalising data. The first approach may be dividing each value by the maximum value within the set. Another approach can be dividing each band by the sum of all bands, or dividing each band of each pixel by the sum for all bands of each pixel [27].

Although noise removal is not essentially a calibration step, it is a very important step to perform before going deep into analysing the data. One important correction to

the raw data that is required is dark current subtraction. As already mentioned, this is a concept were the lens of the scanner is covered with a dark object and an image is taken. The camera should not see any light, but, due to the residual current flowing through the sensors there will be some reading. Hence, dark current subtraction can by calculated by Equation 2.10:

$$DN_\lambda = S_\lambda - D_\lambda \tag{2.10}$$

where DN_λ is the resulting image, S_λ is the raw image, D_λ is the dark current image and λ represents a particular wavelength [27]. Additionally, due to the quantum restrictions of the detector sensors, some spectral bands - usually the ones at the limits of the spectral range - will be noisy. Therefore, it is important to carefully visualise these bands and remove if necessary.

2.4 Summary

This background brings us to the conclusion that for opaque dull surfaces, we can use reflectance spectroscopy technologies to read reflective data. Since the surfaces will be at room temperature, heat radiation can be assumed to be negligible. Therefore, the spectra captured by HSI will be completely representative of the energy level transitions in the visible and close NIR regions (400-1000nm). We also found that the particles that make up the materials have different absorption properties, which will allow us to differentiate between different materials. Moreover, when binders are added, there has to be an adequate PVC ratio such that enough information about the pigment can be captured by the imaging system.

Chapter 3

Literature Review

This chapter contains a review on non-invasive spectroscopy techniques used in this field, and a review on data processing techniques. The aim of this chapter is to evaluate the methods used in recent research and identify possible gaps where improvement to existing methods can be made.

3.1 Imaging Techniques

Imaging spectroscopy techniques are all non-invasive procedures because they do not physically interact with the samples. Most of these spectrometers are very expensive because of their complexity, but they can provide access to sample data which can be revisited at any point in time without requesting additional handling of a painting. Therefore, the increase in the demand for such systems is expected. As mentioned in Chapter 2 these different spectrometers operate on different principles, and some of them measure different data. However, they can also be used as complimentary techniques such that results can be validated. In fact, this was found to be a very popular validation technique in literature, especially when the primary source of identification was Raman Spectroscopy (RS) or X-Ray Fluorescence (XRF).

3.1.1 Raman Spectroscopy

As a form of single-point spectroscopy, RS can be useful when trying to determine the underlying composition of a pigment because it specifically relates with the atomic structure of the pigment. However, when pigments made from similar atomic structures are mixed together, this method may fail as the specific features that relate to the pigments are similar. Moreover, as already mentioned, this method cannot be used on metallic minerals. In the pigment identification field, Appolonia et al. used XRF to

measure surface radiation of atoms and RS to measure molecular features from each layer separately. Generally, both methods are used as complementary techniques to confirm pigment identification results. Also, Daniel et al. have used RS and XRF to confirm results obtained by processing hyperspectral data. Finally, RS has been also used to confirm spectral unmixing results [30]. It is important to mention that although RS is a non-invasive method, and since raman scatter is considered as a weak phenomenon, there is usually the need to take micro samples from the surface and analyse each layer separately as done in [24]. Note that taking micro samples is still considered to be invasive because the sample is being physically breached. Therefore, RS should only be used if other spectroscopy methods are not available, and since it cannot detect metallic compositions prior information about the compositions may be useful.

3.1.2 X-ray Fluorescence

XRF is another spectrometric method that measures the fluorescence that atoms reflect when exposed to X-ray radiation. This method can be used as a spatial-spectral technique, with a good spectral resolution but a low spatial resolution. However, it also has its own limitations when used in the paint restoration field. One of the biggest limitations is that it cannot detect oxides and, as already mentioned in Chapter 2, some pigments are based on oxides. In literature, XRF techniques have been used to determine the thickness of the surface or to identify pigments on a surface. Additionaly, XRF has also been used to validate results obtained by Multispectral Imaging (MSI) [26], and in [31] to determine the types of pigments present on a surface. Brouwer notes that XRF can pose a problem when the sample under examination has a lower density than it should. This is usually the case when pigments are mixed with binders to create different properties. Consequently, the binders decrease the density of the pigment molecules, or affect the Pigment Volume Concentration (PVC) as already seen in Chapter 2.

3.1.3 Fibre Optic Reflectance Spectroscopy

Fibre Optic Reflectance Spectroscopy (FORS) is widely used in the identification of pigments because it makes analysis for single points easy and quick. The information gathered from FORS is spectral information, meaning that no images are acquired. It is usually used as a complimentary method with other spectroscopy techniques such as XRF or RS. In fact, Appolonia et al. used FORS to collect reflective data of the surface, used XRF to collect elemental data of the paintings on the apsidal ceiling and RS to collect molecular data of each layer present on the micro samples. The authors of [24]

conclude that although the results were good at identifying which minerals and pigments were present in the paintings, it was still necessary that FORS, XRF and RS would be used in conjunction to reduce error.

3.1.4 Multispectral Imaging

Another form of spatial-spectral spectroscopy is MSI. Here, a painting is illuminated by natural or artificial light and an image is captured. The information captured by this method is the reflectance that the surface fails to absorb. The use of MSI in identification of pigments has not been frequent. Mostly, because it provides few advantages over other spectrometry techniques. Nonetheless, Deborah et al. used MSI to determine pigments on "The Scream (1893)" which delivered promising pigment classification results, and Cavaleri et al. used MSI to determine the ratios of pigments mixed with a white pigment and XRF to determine the pigments. In [26] results show how MSI indicates a difference in pigment ratios from their reflectance curves. Moreover, their imaging method included the spectral range between 400nm-700nm with a 10nm resolution, totalling up to 30 data points per sample. If we continue to increase the resolution, we can produce a very fine and smooth spectral signature. This introduces Hyperspectral Imaging (HSI), the more advanced form of MSI.

3.1.5 Hyperspectral Imaging

In the above sections, we have seen spectrometric techniques that acquire different types of characteristcs about a sample. For instance, XRF measures the fluorescence intensity, RS measures the raman scatter intensity, and FORS measures a reflectance measure. Since they all measure different metrics, they can be used as complimentary techniques. A similar spectrometric technique to FORS is MSI, except we have seen that in MSI there is the spatial element. By increasing the spectral resolution we obtain HSI. Here, the full advantages of FORS are accompanied by a high spatial resolution. This presents a whole new perspective to analysing spectral data because the algorithms used in RGB images can be extended to HSI. In recent literature, HSI has been applied in the visible range to identify pigments [4], to identify green pigments mixed with linseed oil [28]. Furthermore, Vitorino et al. conclude that red pigments, such as red lakes, are very difficult to be distinguished in the UV-VIS-IR regions [1], while Wu et al. used HSI to identify pigments by going deeper into the short-wave infrared region [29]. Although these works [1, 4, 28, 29, 32] focus on specific pigments from their artwork samples, their results are promising. On the other hand, other literature showed that it is harder to

identify those pigments that had been mixed with linseed oil [28]. At the same time, HSI was also used for the application of segmentation of a painting [4].

Hyperspectral Imaging is a very useful tool which is capable of capturing area-wide spectral information, unlike FORS which can only capture spectral information of a single point. This can provide a better visual of the scene and can eventually provide a better representation of the underlying features. A good example is the work of Attas et al. where underlying sketches were visualised using HSI [3].

3.2 Pigment Processing Techniques

The data captured by the HSI system is multi-dimensional, where each dimension is a specific wavelength. Since the spatial resolution is also of high quality, data processing techniques can be applied to extract the desired information. This information can either be spectral features, or spectral features linked with spatial features. Before any of this can be done, the acquired data must first be cleaned from any noise that is acquired. Since the system is digital and sensors are used to convert analogue data to digital data, small fluctuations between neighbouring pixels may be visible. In this section, a review of existing methods on data filtering, feature extraction, and data classification is provided.

3.2.1 Filtering

In signal acquisition, different equipment may result in different signatures due to the effect of random noise. Filtering the signals mitigates this effect and helps to produce a smooth signal that becomes equivalent as to when captured with less noisy equipment. Since we are dealing with hyperspectral data, it is ideal to filter the data both spectrally and spatially. Spectral noise may occur from the transitioning of bands, whereas spatial noise may occur from the different sensitivities of neighbouring sensors. In HSI, due to the physical limitation of the sensors, low SNR will be often found at the edges of the bandwidth. In fact, in most literature, these wavelengths are truncated [4, 5, 31, 33]. At the same time, spectrally filter their multispectral data cube such that a smoother transition between bands could be obtained [26].

3.2.2 Feature Extraction

Image features can be several different things. In most cases, these are what humans can interpret, however, sometimes we can also extract features hidden to the human eye. Examples of features are for example; edges, colours, objects and textures [34].

Edge features are useful to detect overlapping features having different contrasts and can be detected by algorithms such as, Sobel [35] and Combination of Receptive Fields (CORF) [36]. In paintings, edge detection can be useful to detect prominent edges created by adjacent pigments which have spectrally different signatures. There is often the case, however, in which two paints having similar spectral signatures are placed next to each other to create the illusion of a smooth transition. Based on our experiments, this smooth change here will not be able to be captured by an edge detection method. Similarly, object detection can be very useful to detect different objects in a painting, however, without the possibility of identifying the type of paint. For textures, it is possible to separate areas by looking at the different properties of textures [37]. In paint, this can translate to the different brushstrokes one can leave when applying the paint. Unfortunately, however, this does not relate to separating different paints, since the texture depends on the brush and artist rather than the paint itself. Colour features are used to distinguish between areas by analysing the pixel intensities and can be obtained by using algorithms such as the K-means [37]. Since different paints radiate light at different energy levels, this method can, by far, be the most useful for the separation of paint.

Recent literature in the analysis of paintings show that most of the paint identification is performed in a pixel-by-pixel manner. This approach is easy and time efficient to perform, but at the cost of having inaccurate identifications, especially at paint borders. Such approaches can be seen in [1, 5, 38]. Other approaches that perform some feature extraction prior to the identification are less time efficient, but they can group pixels together based on their spectral relationships. Such methods, are for example, clustering techniques. Clustering-based feature extraction has also been applied to HSI. In fact, in [28], a Density Based Spatial Clustering of Applications with Noise (DBSCAN) method was applied to mock-up samples created by the authors. The authors first reduced the dimensionality of their HSI data by applying Barnes-Hut t-Distributed Stochastic Neighbour Embedding (BH-TSNE), which maps multi-dimensional data to, but not limited to, two-dimensional data. Here, DBSCAN is more suitable than other clustering methods such as K-means, since K-means clusters are restricted to have a spherical or convex shape, whereas in DBSCAN the shapes of clusters may be arbitrary. However, the results of this paper do not solve the fragmentation problem as seen in Figure 3.1(a). Other techniques that also include spatial features, like the Watershed Technique (WT), have also been extended to HSI [39]. Similarly to the original WT, a gradient image is created to define the 'ridges' per wavelength image. These images are summed together and a threshold then separates 'true' boundaries from

'false' boundaries. Although this method can be very accurate at defining boundaries, it is completely dependent on the defining power of a 'ridge' which is accumulated per wavelength. Therefore, two paints having similar spectral intensities (such as two reds) will be difficult to separate because of the small change in gradient in every wavelength image. Other techniques, such as Simple Linear Iterative Clustering (SLIC) [4] have also been extended to HSI. The authors extended the SLIC algorithm to cluster the image by processing hyperspectral data. In comparison to the DBSCAN approach, the resulting identification is not fragmented. However, an over-segmentation problem emerges as shown in Figure 3.1(b). To better understand the over-segmentation problem of SLIC, a more detailed explanation of the algorithm is given in Section 3.2.2.1.

In the RGB domain, region merging techniques have already been applied to the results obtained by algorithms such as SLIC [40]. Wu et al. compare the resulting segments of SLIC in terms of standard deviation, and decide whether they should be merged or not. The results show that the resulting regions are consistent with the human perception of the used images. Similarly, in [41], region merging is applied on grayscale images for segmentation results from SLIC0, a variant of SLIC which allows a dynamic compactness feature. Depending on the starting value of superpixels, region merging based techniques show that the over-segmentation problem can be reduced by a large amount. Another region merging method applied in the RGB domain includes the separation of the foreground from the background through an interactive approach using a brush to obtain texture samples from the image. In this way, the algorithm can identify regions and merge them by looking at these samples [42]. In [43], the authors combine the advantages of over-segmenting an RGB image with uniform superpixels and the advantages of finding edges of an image. By overlapping these two results, Li et al. were able to modify the superpixels that contain an edge within them such that the superpixels can strongly adhere to boundaries. However, superpixels that do not contain an edge, remain unchanged, implying that the homogeneity property of these superpixels is not optimised. Another SLIC-based region merging approach is proposed in [44]. The method combines the resulting superpixels from SLIC, and the proposed 5-D Spectral clustering and boundary-focused region merging stages. The algorithm also includes adaptive thresholds. Results show that the algorithm can produce large segments, however, the homogeneity property is not clearly defined. The method appears to focus on perceptual grouping of objects. This grouping may be problematic when dealing with segmentation for paintings for the purpose of paint identification, since regions containing different paints need to be segmented even if they lie within the same perceptual group. From this literature, we can conclude that the SLIC has been very

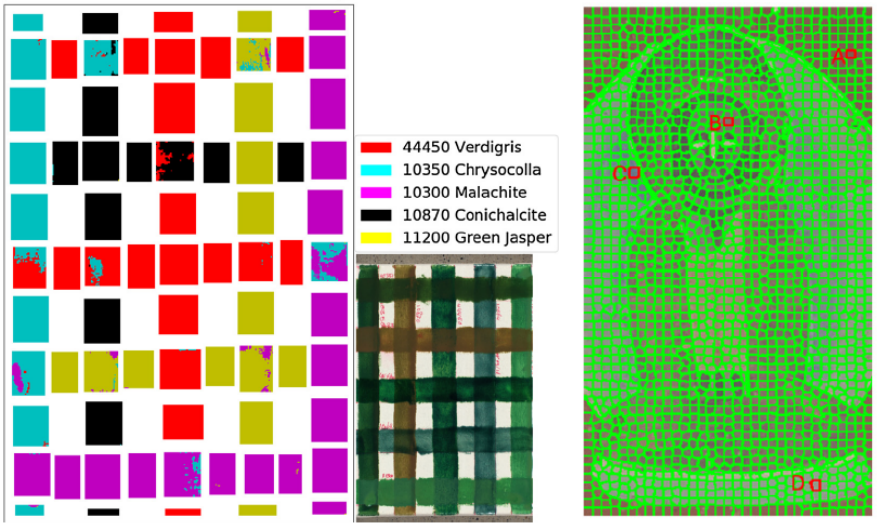

(a) (b)

FIGURE 3.1: Comparing two different approaches to pigment analysis. In (a) a pixel-by-pixel approach is developed which results in fragmentation of pixels [28]. In (b) a superpixel approach is addressed which eliminates fragmentation but introduces over-segmentation [4].

widely applied to region merging approaches. Hence, we can study the SLIC superpixel approach for HSI, taking special care to separate paints that have a very similar hue, but have a different recipe.

3.2.2.1 Simple Linear Iterative Clustering

The SLIC [45] algorithm is, as the name suggests, a form of clustering algorithm, in which the final result contains superpixels. It incorporates a slight variation of the K-means algorithm to create the superpixels. Achanta et al. compared results from the algorithm to other state-of-the-art superpixel segmentation algorithms which show that SLIC outperforms most of the mentioned algorithms in both speed and boundary adherence [45].

SLIC requires the number of superpixels k, and a compactness factor m which regulates the shape of the superpixels. A number of superpixels is predefined in a grid like manner, of size $S \times S$ where S is equal to $\sqrt{\frac{N}{K}}$, with N being the total number of pixels in the image. Then the algorithm calculates the distance between each pixel that falls within the search region of each cluster using Equation 3.1 and Equation 3.2. The search region is that of $2S \times 2S$. If the distance D is smaller than the previous saved distance, the distance and label are replaced. Once all clusters have been checked for the pixels that fall within their search region, the centroids are updated and the algorithm

starts again until a termination criterion is met. For an RGB image, the feature vector mentioned will consist of five values, three values for red, green and blue, and two values for spatial location. In this case the Euclidean distance is used as a distance metric:

$$d_c = \sqrt{(l_j - l_i)^2 + (a_j - a_i)^2 + (b_j - b_i)^2} \tag{3.1}$$

and,

$$d_s = \sqrt{(x_j - x_i)^2 + (y_j - y_i)^2} \tag{3.2}$$

where d_c and d_s are the resultant distances for colour and location respectively, l, a and b represent the colours in CIELAB, while x and y represent the spatial location. Subscripts j and i represent pixel and cluster respectively. Euclidean Distance (ED), then, can also be used to calculate the total combined distance as shown in Equation 3.3:

$$D = \sqrt{\left(\frac{d_c}{m}\right)^2 + \left(\frac{d_s}{S}\right)^2} \tag{3.3}$$

where D is the total distance, m is the colour constant and S is the maximum distance from the cluster [45]. In Equation 3.3, S is said to be the maximum spatial distance that could be achieved for each region, where S is equal to $\sqrt{\frac{N}{K}}$. It is used to normalise the spatial distance, while m on the other hand is used to place a weighting on the colour distance. Achanta et al. claims that it would be quite difficult to find a predefined value to normalise the colour distance since each cluster is different than the other, and that images differ from each other. This way m also works as a compactness factor, where if closer to 0 the superpixels will be less regular in shape.

Li and Wan made use of the SLIC algorithm to segment their hyperspectral data. They made changes to Equation 3.1, which became Equation 3.4 such that it works for hyperspectral data:

$$D = \sqrt{\sum_{\lambda} (\mathbf{r}_{j\lambda} - \mathbf{r}_{i\lambda})^2} \tag{3.4}$$

where $\mathbf{r}$ is the wavelength vector. This variation of the SLIC can be used for paintings with distinct pigments [4], but due to the nature of ED, it will also separate patches with different brightness even if they are made of the same pigment. In this study the aim is to mitigate this problem without affecting the boundary adherence of the algorithm.

3.2.3 Pigment Identification

The complexity of identifying paints depends on the painting being analysed. An old painting can simplify the complexity of identifying the paints by finding which artist painted the artwork and the place of origin and date, because different cultures used different and specific selection of pigments. However, it is not always that easy. Moreover, as knowledge regarding pigments became more globalised the selection of pigments became more vast, hence, the complexity of identifying paints increased. As mentioned in Chapter 2 there are pigments that have been used in their pure form, and there are pigments that have been mixed with other pigments to create specific shades of colours.

For the former, classification techniques are usually accurate enough to identify pure paints, with some exceptions depending on factors such as the range of wavelength. In literature for paint identification, Spectral Angle Mapper (SAM) is very commonly used as the metric that compares the paints from the sample to that in a library [4, 28, 29]. It compares the spectral angle between two spectral signatures. A variant of SAM is Spectral Correlation Mapper (SCM), which removes the mean from the spectral signatures to focus on shape features rather than intensity. It has been used in some literature to compare the results against SAM [5, 38, 46]. Although these two methods are distance metric techniques, they classify the pigments by looking for the minimum or maximum value when compared to the other pigments in the library. Although both SAM and SCM are widely used, they have limitations. In SAM, any spectral variability in intensity will cause inaccuracies, whereas any spectral variability in wavelength will cause inaccuracies for SCM. In [47], the authors modified SAM to Kernel Spectral Angle Mapper (KSAM) by making each term of the SAM be a result of a kernel function, as will be explained in more detail in the following sections. The modification increased the overall accuracy of the conventional SAM by 3% [47]. In other research [32] a 1-dimensional Convolutional Neural Network (CNN) has been trained to make pixel-by-pixel classification of pigments on artworks.

On the other hand, when two pigments are mixed together, conventional methods such as SAM and any of its variants, cannot be used because the method will not be capable of identifying that there is a mixture in the spectrum. Also, since a mixture of paints can be considered to be a linear superposition model, the addition of two signatures may result to be more similar to another pigment which is not present in the mixture but is available in the library. Therefore, to identify pigments in mixtures, spectral unmixing techniques are required. Spectral unmixing techniques can be linear and non-linear. Strictly speaking, because light can only be considered to be linear in vacuum, non-linear

models should be used for best accuracy. However, these methods require a lot of prior knowledge on the mixtures, such as particle size, mass and molecular structure [11]. In some literature, non-linear models have been used to identify pigments in mixtures [31, 48]. Since the prior knowledge regarding the pigments is not always something one has at hand, most literature choose to use linear models to identify pigments in mixtures. In fact, linear models such as Non-negative Matrix Factorisation (NNMF) and Independent Component Analysis (ICA) have been used [18]. Additionally, in [18] the abundance ratio of the involved pigments in mixtures have been estimated and the authors have shown how their method improves on the accuracy of existing methods. A limitation of the work in [18] is that they use the prior knowledge of knowing a priori, which paints are in the mixture in order to successfully estimate the abundance ratios. In the case of non-linear models, neural networks are more commonly used [31, 48], however, an application of the Kubelka-Munk (KM) theory was also found in [20]. For the former, neural networks have been introduced to many areas of research. Their limitation is that a lot of data is required for a network to learn. However, it could be very advantageous because with one network, a whole spectra of mixtures can be identified. The problem here is that one has to create all of the possible mixtures for the neural network to learn. In [48] a deep neural network was trained to identify mixtures of watercolours. In the case of the KM theory [20], a model for estimating the colours in an RGB image was proposed. The model is based on existing dataset containing the absorption and scattering coefficients of the known pigments in the RGB image. From these known coefficients, the pigments and their concentration were acquired for each pixel in the image. The authors state that the proposed method is similar to the NNMF due to its convexity, but is non-linear unlike NNMF because it incorporates the KM theory. The non-linearity increases the accuracy of extracting reliable absorption and scattering coefficients. To simplify the complexity of computations, the substrate and thickness were assumed to be equal to one. Further assumptions with regards to the image were also taken [20].

In the following subsections, we will be going a bit deeper into the mentioned algorithms for paint identification. First we will cover the pure paint classification techniques and then we will cover spectral unmixing.

3.2.3.1 Spectral Angle Mapper and Variants

Originally created by NASA for classification of endmembers in remote sensing applications, SAM is widely used in pigment identification. It measures the total angle difference between two endmembers for each band in the spectrum, see Figure 3.2,

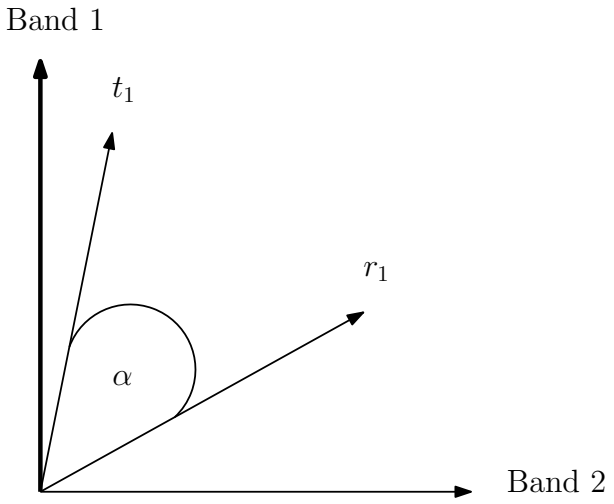

FIGURE 3.2: Illustration of how SAM calculates a similarity measure

in which one endmember is the sample and the other is any from a reference library. Equation 3.5 is the formula for measuring the angle between two signatures:

$$\alpha = \cos^{-1} \frac{\sum_{i=1}^{nb} t_i \cdot r_i}{\sqrt{\sum_{i=1}^{nb} t_i^2} \sqrt{\sum_{i=1}^{nb} r_i^2}} \tag{3.5}$$

where $\mathbf{t}$ is the sample endmember, $\mathbf{r}$ is the reference endmember and nb is the total number of spectral bands. Essentially, this follows the inner product of two vectors with the inverse cosine transforming the result into an angle. The performance of the algorithm has been confirmed by various studies [4, 5, 38]. However, most of the time it is outperformed by its variant, the SCM [49]. The main difference between the two is that SCM removes the mean from each endmember and uses the Pearsonian Coefficient [49], which increases the power of a match. From Equation 3.5, the formula for SCM is given by Equation 3.6:

$$c = \frac{\sum_{i=1}^{nb} (t_i - \bar{t}) \cdot (r_i - \bar{r})}{\sqrt{\sum_{i=1}^{nb} (t_i - \bar{t})^2} \sqrt{\sum_{i=1}^{nb} (r_i - \bar{r})^2}} \tag{3.6}$$

where $\bar{t}$ is the mean of the sample signature and $\bar{r}$ is the mean of the reference signature. Various aforementioned studies report how the SCM outperforms SAM in the identification of pigments. The only drawback of SCM is that for a slight shift in wavelength between two endmembers the performance will degrade.

Liu and Yang have modified the SAM to be a kernel based algorithm. It still measures the angle between two signatures but instead of measuring the spectral angle directly, a kernel is used to measure the angle [47]. The paper also provides a table of different

functions used for the kernels such as Sigmoid, Radial Basis Function and Polynomial. Modifying Equation 3.5, the kernel based formula becomes:

$$\alpha = \cos^{-1} \frac{\boldsymbol{K}(t_i, r_k)}{\boldsymbol{K}(t_i, t_m) \cdot \boldsymbol{K}(r_i, r_k)} \tag{3.7}$$

where Equation 3.8 expands one term of $\mathbf{K}$ from Equation 3.7 to show the kernel function:

$$\mathbf{K}(t_i, t_m) = \mathbf{\Phi}(t_i)^T \cdot \mathbf{\Phi}(t_m) \tag{3.8}$$

and $\mathbf{\Phi}$ is the applied function, which expands Equation 3.7 into Equation 3.9:

$$\alpha = \cos^{-1} \frac{\sum_{i=1}^{nb} \mathbf{\Phi}(t_i) \cdot \mathbf{\Phi}(r_i)}{\sqrt{\sum_{i=1}^{nb} (\mathbf{\Phi}(t_i))^2} \sqrt{\sum_{i=1}^{nb} (\mathbf{\Phi}(r_i))^2}} \tag{3.9}$$

According to the results in [47] the KSAM, when compared to SAM, can improve the overall accuracy by up to 3.5%.

3.2.3.2 Spectral Unmixing

Spectral unmixing could be based on linear or non-linear methods. For linear methods, information such as knowing the pure endmembers, or knowing boundary endmembers that fit a minimum volume simplex is beneficial. However, one can also perform linear unmixing without prior information, which results to be less accurate. On the other hand, nonlinear methods require much more information to be able to find endmembers, such as particle size of the material. There are various studies that have already implemented such techniques to determine the different pigments in mixed regions [5, 31, 38], all showing promising results. In general, by the end of an unmixing stage, a set of endmember spectra present in the sample region should be obtained. In this study, the focus is on linear methods, as physical properties regarding our pigments are not within our reach.

The Simple Identification Split Augmented Lagragian (SISAL), is a linear unmixing algorithm that is part of the minimum volume simplex class. Such algorithms are: Minimum Volume Enclosing Simplex (MVES), Minimum Volume Simplex Analysis (MVSA), and Non-negative Matrix Factorisation - Minimum Volume Transform (NMF-MVT) [50]. The SISAL solves a hard nonconvex problem by using subnormal convex constraints [50]. Grabowski et al. used the algorithm to identify the endmembers in hyperspectral data of an entire painting; a template of different green pigments. Their results show that SISAL can estimate a good simplex volume for pigments that are made up of multiple endmembers.

Other statistical methods used for the spectral unmixing of paint regions include ICA and Principal Component Analysis (PCA). The goal of ICA is to obtain independent components which would correspond to the individual paint vector as the basis vector. By assuming independence of the components the ICA can perform blind source separation of mixed signals without requiring any assumptions on the Gaussianity of the components. On the other hand PCA assumes that this data can be segmented into uncorrelated and orthogonal basis. It is therefore more suited for compressing the feature vector representing the data [3, 29]. In [18], a variant of ICA called FastICA has been applied to spectral unmixing of paint mixtures, in which the algorithm was outperformed by another linear matrix factorisation algorithm called NNMF. The goal of [18] was to find the most suited algorithm for extracting the true endmembers in their mixtures.

The NNMF method is a linear unmixing method that factorises a matrix into two other matrices [18], with the only constraint being that all three matrices must not have any negative elements:

$$\mathbf{V} = \mathbf{WH} \tag{3.10}$$

where $\mathbf{V}$ is the spectrum matrix, $\mathbf{W}$ is the basis matrix and $\mathbf{H}$ is the abundance matrix. Each row of the basis matrix represents each spectrum separately, while the abundance matrix consists of values that represent the ratios of the pigments. This is only possible by having different update rules for each matrix. The update rule for $\mathbf{H}$ is given by Equation 3.11 and that for $\mathbf{W}$ is given by Equation 3.12:

$$\mathbf{H}_{t+1} \longleftarrow \mathbf{H}_t \cdot \frac{(\mathbf{W^T V})_t}{(\mathbf{W^T W H})_t} \tag{3.11}$$

$$\mathbf{W}_{t+1} \longleftarrow \mathbf{W}_t \cdot \frac{(\mathbf{V H}^T)_t}{(\mathbf{W H H}^T)_t} \tag{3.12}$$

where T denotes the transpose of the matrix and t denotes the iteration. In HSI, Shuqiang et al. used NNMF to compare its results with those of the aforementioned ICA algorithm. According to the results in [18], NNMF is the better option in the identification of pigments. This is encouraging to continue analysing this algorithm for ratio estimation of various pigments in a single mixture. Even though the application was not in the restoration of artworks, but rather regarding musical instruments, Dessein et al. have made use of NNMF in a different way. They first learned a basis matrix $\mathbf{W}$ by using true signals, which could be pigments from our library. Then live signals are fed and an abundance matrix $\mathbf{H}$, is learned based on the pre-learned basis matrix $\mathbf{W}$. Results from [51] are very promising, and in this study, an attempt is made to implement this

TABLE 3.1: Summary of literature focusing on the type of equipment and methods used.

Year	Author/s	Equipment					Spectral Range (nm)	Method	Classification		
		XRF	RS	FORS	MSI	HSI			DF	SU	NN
2013	Cavaleri et al. [26]	✓			✓		400-700	Cross-sectional			
2015	Vitorino et al. [1]					✓	400-900	Point-based			
2016	Rohani et al. [38]					✓	400-900	Point-based		✓	
2017	Wu et al. [29]					✓	1000-2500	Point-based	✓		
2018	Grabowski et al. [28]					✓	1000-2500	Clustering-based	✓		✓
2018	Rohani et al. [31]	✓			✓		400-900	Point-based			✓
2018	Daveri et al. [8]	✓	✓			✓	2700-5500	Clustering-based	✓		
2018	Li and Wan [4]					✓	400-1000	Superpixel-based	✓		
2020	Kleynhans et al. [32]					✓	400-950	Point-based			✓
2020	Shuqiang et al. [18]					✓	350-2500	Point-based	✓	✓	

DF Distance Function
SU Spectral Unmixing
NN Neural Networks

unmixing analysis to identify different pigments in mixtures.

In [18], the NNMF method is based on the reflectance of the samples. In this work, an attempt is made to apply NNMF to absorption and scattering coefficients obtained by applying the KM Theory in [20], instead of reflectance. Since NNMF is a linear unmixing model, and reflectance is nonlinear, by separating the reflectance into absorption and scattering coefficients a more linear signature should be obtained. Therefore, a significant improvement over the accuracy of paint identification with NNMF should be observed.

3.3 Summary

From this section we can conclude that it is much more advantageous to use non-invasive spectroscopy over invasive methods. Non-invasive techniques can improve the quality of data analysis, as well as provide the restorers with reusable data. Moreover, by using HSI, artwork exposure to radiation is limited, whereas in XRF artworks have to be exposed to unnecessary harmful radiation. Therefore, in this study, the data will be acquired using a hyperspectral camera. In Table 3.1, a summary regarding important points, such as the method of acquisition, data acquisition methods and data processing methods, is illustrated. Based on the review performed with regards to data processing techniques used in the art restoration field, we come to the following conclusions:

- We can perform feature extraction to separate homogeneous regions into different segments while minimising over-segmentation and keeping a high adherence to pigment boundaries.

- We can improve on the already existing variants of SAM to further increase the accuracy of identifying true pigments.
- We investigate accuracy measures for mixtures using NNMF for the KM absorption and scattering coefficients method.

Since in [1] it was concluded that red paints are generally difficult to differentiate in the visible region, we include three red paints in our selection, alongside two green paints, three blue paints and one yellow paint. This selection should provide enough variety in the spectral signatures to provide an informative study. These paints should also show the robustness of our methods when compared to others.

Chapter 4

Methodology

In this chapter, we explain the equipment and materials used for data acquisition and analysis. Moreover, we propose a methodology to segment the hyperspectral data into regions having the same spectral properties, and identify the paint within these regions taking extra measures if a mixture of paints is detected.

4.1 Materials and Equipment

Paint can be expressed as pure pigment, a mixture of pigments, and even layered pigments. The pure pigment case is fairly simple to analyse because the reflectance being measured can be directly matched to an existing library. In mixture of pigments, a homogeneous mixture of the elements can be analysed if the total number of different pigments in the mixture is known, such as by using Non-negative Matrix Factorisation (NNMF) or Kubelka-Munk (KM) Theory. When pigments are layered, the identification becomes more complex because various factors affect the reflectance signature. For example, two pigments, A and B, have different opacities, in which A is more opaque than B. Therefore, if A is layered over B, more of A will be reflected or absorbed and, hence, less light will reach layer B. Conversely, when B is layered on top of A, because it is less opaque, more light will reach A. This means that the final reflectance values, for the same thickness levels, will vary depending on which pigment is on top. In fact, it has been observed that layered paint can pose some level of difficulty in being separated; in [28] the layered sections are fragmented which shows a level of uncertainty, that is a region made of pigment A layered over B is sometimes classified as A and sometimes as B, as shown in Figure 3.1(a). We argue that it would be ideal if the entire region is classified as both A and B, since both pigments are present. In this work, therefore, we

TABLE 4.1: The mass to mass paint mixing ratios for the homogeneously paint mixture samples with Paint 1 (50%) and Paint 2 (50%).

Mixture	Paint 1 (mg)	Paint 2 (mg)
Cadmium Red Hue + Viridian Hue	840	837
Cadmium Red Hue + Cerulean Blue Hue	822	821
Cerulean Blue Hue + Viridian Hue	856	853
Lemon Yellow Hue + Viridian Hue	327	308
Lemon Yellow Hue + Cadmium Red Hue	454	509
Lemon Yellow Hue + Cerulean Blue Hue	543	592
Emerald Green + Lemon Yellow Hue	711	716
Emerald Green + Permanent Rose	452	463
Emerald Green + French Ultramarine	307	303
Phthalo Blue + Permanent Rose	379	418
Phthalo Blue + Lemon Yellow Hue	207	181
Permanent Rose + Lemon Yellow Hue	318	323

prepare mock-up samples to try and imitate the above problems, and to test whether our methods work under different problems.

The pigments used are pre-mixed with linseed oil from Winsor & Newton ©: Viridian Hue, Emerald Green, French Ultramarine, Cerulean Blue Hue, Phthalo Blue, Permanent Rose, Permanent Geranium Lake, Cadmium Red Hue and Lemon Yellow Hue. The mock-up samples created for this work included three samples for each mixture, specifically ground truth sample, two samples for analysis and one sample for testing the final algorithm. The ground truth sample includes patches of all pigments. The mock-up samples for analysis follow a grid like pattern - creating the layered patches - made up from the first six pigments from Winsor & Newton. All samples are prepared on stretched canvas primed with Titanium Dioxide. A second sample for analysis include irregular shaped areas of oil-pigments of different ratios for the nine mentioned pigments. The testing sample follows an irregular pattern to simulate a painting made up of all pigments. Some homogeneous mixtures of approximately 50:50 mass to mass ratio were also created and applied to a swatch-like painting sample. The mass-to-mass ratios of the mixtures can be noted in Table 4.1.

For validation purposes and robustness testing, a similar mixing procedure with varying mixing ratios has been carried out for the same mixtures. The mixing ratios are set as 60:40 and 40:60 and are noted in Table 4.2 and Table 4.3, respectively.

It is important to note that, even though not within the scope of this work, certain paints may contain multiple elements, sometimes used for other paints. Moreover, the elements used to create these paints may vary over the years, usually depending on health

TABLE 4.2: The mass to mass paint mixing ratios for the homogeneously paint mixture samples with Paint 1 (60%) and Paint 2 (40%).

Mixture	Paint 1 (mg)	Paint 2 (mg)
Cadmium Red Hue + Viridian Hue	342	228
Cadmium Red Hue + Cerulean Blue Hue	472	315
Cerulean Blue Hue + Viridian Hue	423	282
Lemon Yellow Hue + Viridian Hue	310	206
Lemon Yellow Hue + Cadmium Red Hue	358	238
Lemon Yellow Hue + Cerulean Blue Hue	490	310
Emerald Green + Lemon Yellow Hue	334	216
Emerald Green + Permanent Rose	400	270
Emerald Green + French Ultramarine	300	200
Phthalo Blue + Permanent Rose	296	200
Phthalo Blue + Lemon Yellow Hue	268	179
Permanent Rose + Lemon Yellow Hue	442	295

TABLE 4.3: The mass to mass paint mixing ratios for the homogeneously paint mixture samples with Paint 1 (40%) and Paint 2 (60%).

Mixture	Paint 1 (mg)	Paint 2 (mg)
Cadmium Red Hue + Viridian Hue	124	186
Cadmium Red Hue + Cerulean Blue Hue	193	290
Cerulean Blue Hue + Viridian Hue	158	238
Lemon Yellow Hue + Viridian Hue	150	225
Lemon Yellow Hue + Cadmium Red Hue	176	262
Lemon Yellow Hue + Cerulean Blue Hue	206	302
Emerald Green + Lemon Yellow Hue	166	234
Emerald Green + Permanent Rose	295	442
Emerald Green + French Ultramarine	381	571
Phthalo Blue + Permanent Rose	263	385
Phthalo Blue + Lemon Yellow Hue	254	381
Permanent Rose + Lemon Yellow Hue	151	227

and safety guidelines [52]. In Table 4.4, we list down the different elements used in each of the paints in our palette. The information regarding the elements is acquired from the back of the tube, which lists all elements used for that paint. In the case of Winsor & Newton, the paint code is referred to by Colour Index Generic Names. For example, Cobalt Blue is referred to as Pigment Blue 28, or abbreviated to PB28 [52]. From Table 4.4, we can observe that some paints contain multiple pigments such as Cerulean Blue Hue, which contains some Phthalo Blue.

TABLE 4.4: Noting the different elements used to create the different paints

Paint Name	Pigment Code/s
Lemon Yellow Hue	PY3
Viridian Hue	PG7
Emerald Green	PG36, PY74, PW5
Cerulean Blue Hue	PB15, PG7, PW4
French Ultramarine	PB29
Phthalo Blue	PB15
Permanent Geranium Lake	PR209
Cadmium Red Hue	PR118, PR170
Permanent Rose	PV19

The SPECIM FX10e hypspectral imager is used to capture the data samples. This hyperspectral camera has a spectral range of 400-1000nm with a spectral FWHM of 5.5nm resulting in 448 wavelength channels and spatial resolution of 1024 pixels. It has a very high Signal to Noise Ratio (SNR) 600:1, a blocking order filter and an electromechanical shutter that deals with dark backgrounds. Specifically, the blocking order filter reduces the effect of multiple wavelengths falling on the sensors [53].

4.2 Data Acquisition

The camera sensor reads intensity values radiated from the sample surface. The software that comes with the camera also returns intensity values. The SPECIM FX10e imager saves the raw data as a 1-D vector, in which the correct dimensions for 3-D conversion (pixel width, pixel height, channel wavelength) can be found in a header file. We can use this information to reconstruct the image data cube from the row vector provided. Therefore, the first step after acquiring the data is to reconstruct the raw 1-D vector into a 3-D data cube. Once the image cube is reconstructed, we must make sure that the illumination on the surface is uniform. To check for the uniformity of illumination on the surface of the sample, a spectralon (usually 99% reflective surface) block accompanies the sample when the data is collected. Since the hyperspectral imager is a line scanner - or pushbroom - the variance in illumination over the surface is assumed to be only in the direction that is stationary. That is, if direction x is the direction the table is transvering, then y is the stationary axis. This variation can be reduced by using the white and dark references obtained during data acquisition to do background subtraction.

Background subtraction is required to remove background noise from the sample acquired, see Equation 2.9 in Section 2.3.5.2. To account for the variation in illumination,

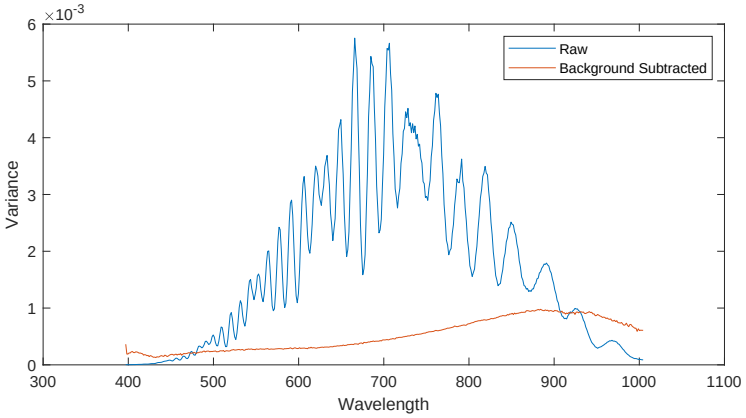

FIGURE 4.1: Variance of a line scan against wavelength to show the ambient correction by background subtraction. Varying illumination is also normalised.

the white reference is reshaped to a 2-D matrix instead of 1-D. That is, taking the mean illumination on the white sample across the direction the table is moving, instead of both directions. Had the illumination been uniform across the white reference and because it is a spectralon surface, then it could be assumed that each point on the surface has the same amount of incident light. This would allow us to reduce the white matrix to a single value per wavelength, however, it is not the case.

After performing the background subtraction, we confirmed the reduction in illumination variance by plotting the difference in illumination across the y direction, as shown in Figure 4.1. Moreover, it is important to show that this stage does not change the paint features, but rather removes the background noise. In Figure 4.2 a plot relating to the change in data between the raw and background subtracted data is illustrated. From this plot, one can conclude that although background subtraction causes the ambient features to change, it does change the relationship between the features. One of the reasons why this happens is because the camera itself is a line scanner, meaning that variance in illumination will only happen in one direction whereas the random noise from the camera sensors will happen across all the image. The next section specifically focuses on the method of filtering the noise caused by the data acquisition equipment.

Moreover, we normalise the data by the maximum value acquired by the spectralon block that is beside the mock-up during acquisition. In this way, we normalise separate mock-up samples, even if they were exposed to different light intensities.

For comparative reasons, discussed in Chapter 5, we transform the Hyperspectral Imaging (HSI) data to RGB data. The conversion to the RGB domain was performed by taking the three dominant wavelengths for each colour; Blue (450nm), Green (560nm) and Red (690nm) which are then concatenated together to create a single RGB image.

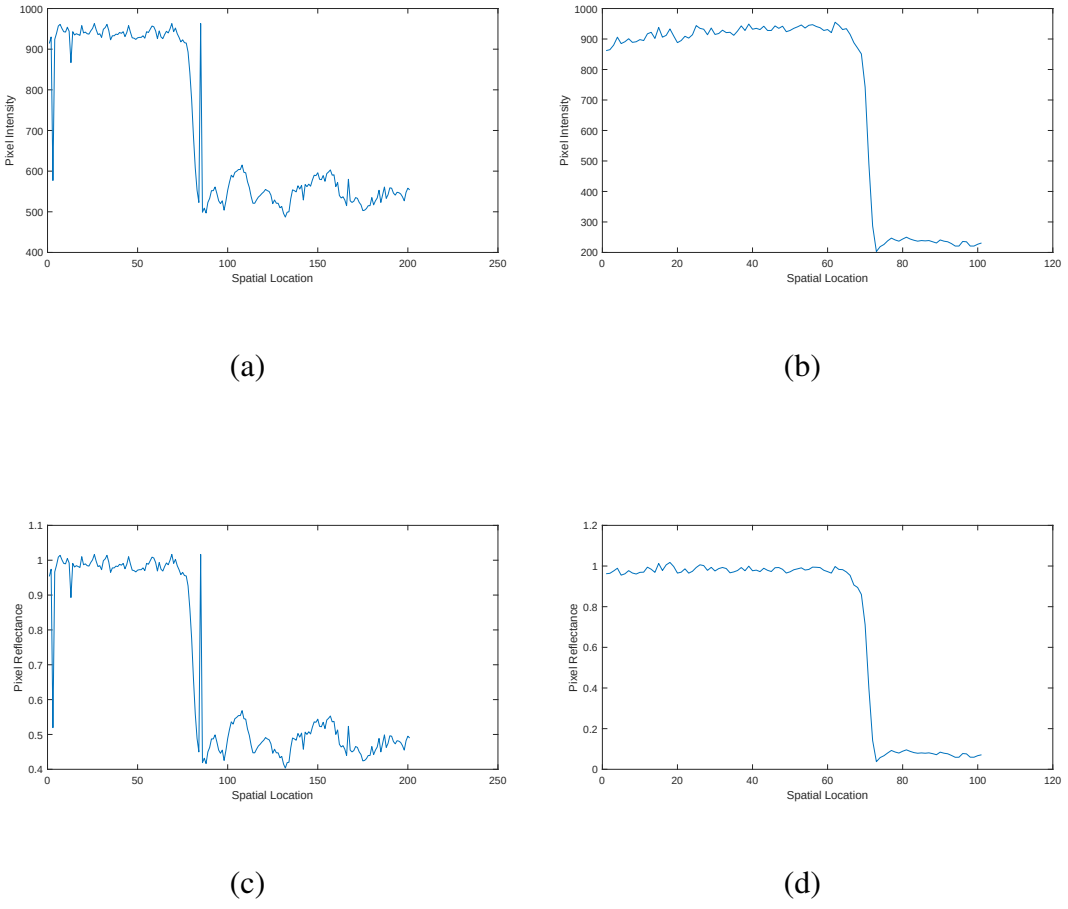

FIGURE 4.2: This illustration shows the effect of the background subtraction applied to an image taken at a single wavelength. Plots (a) and (b) show a horizontal and vertical scan-line from this image giving the raw pixel intensities obtained from the hyperspectral camera. Plots (c) and (d) show the result of background subtraction which removes ambient noise and converts the raw pixel intensities into reflectance values. These plots demonstrate that spectral features are retained after background subtraction.

4.3 Filtering

In imaging, the most common type of noise is when neighbouring pixels vary in intensity due to differences in sensor sensitivities. In some cases, especially in older equipment, sensors may become faulty causing large discrepancies between neighbouring pixels. In grayscale images, this type of noise is known as salt and pepper. Therefore, filtering techniques are required to improve the result of a bad pixel. Bartyzel mentions some filtering methods; moving average filter, median filter, and Kuwahara filter. In the following subsections, an explanation on the type of noise acquired and the necessary steps taken to mitigate the effects are explained.

4.3.1 Spatial Domain Filtering

One type of noise present in our data that may interfere with analysis is that of misplaced lines; this can also be referred to as data corruption. Although the cause of this noise is known - a communication limitation - it cannot be fixed unless the system is upgraded. Therefore, we aim to mitigate this problem by using filtering methods, without losing important information, such as edges.

The Median Filter finds the value at the middle of a sorted vector in ascending order [54]. The adaptive median filter is an improvement to the median filter. It allows a dynamic window size that changes according to the values in the sorted vector. It works by initially sorting the vector in ascending order, saving the minimum, maximum and median values. If the median value equals either the maximum or the minimum, the window size increments by 1 and the sorting is reapplied. Otherwise, the median value is used as the new data. Although this method improves noise reduction, it can be computationally heavy depending on the size of the image [54].

Bartyzel argues that filtering methods such as moving average filter suffer from some degree of blurriness. Although blur is obtained when reducing noise, it is also possible to mitigate this effect, to preserve sharpness on important edges. The Kuwahara filter is an edge-preserving filter. It works by splitting the window into four different areas of equal size; each area includes the test pixel. The variance and average of all areas are calculated and the output pixel will take the average respective to the area with smallest variance. Although this filtering method returns promising filtering results, improvements can be made at low SNR values because more information regarding neighbouring features can be gathered. Bartyzel proposed a new algorithm that combines the adaptive median filter and Kuwahara filter to improve on the results of the Kuwahara filter. An illustration of how the four different quadrants can grow adaptively and independently from each other is shown in Figure 4.3.

The adaptive Kuwahara filter [54] aims to reduce blur and pixellation at higher window sizes. Although it follows the adaptive median filter, it does not incorporate the median aspect but the adaptivity. It works by initially splitting the window into four quadrants as done for the Kuwahara filter. Each area starts with four pixels (including the test pixel) and in an iterative manner, the variance and average of each area are calculated and saved. Iteratively, each area increases its size by one pixel in each dimension, so from four pixels it will now be nine pixels (including the test pixel) under a single quadrant. Once again, the variance and average are calculated. If the variance of an area is larger than the variance of the same area in the previous iteration, then the current area size is discarded and this area stops growing. If, on the other hand, the variance of an area is

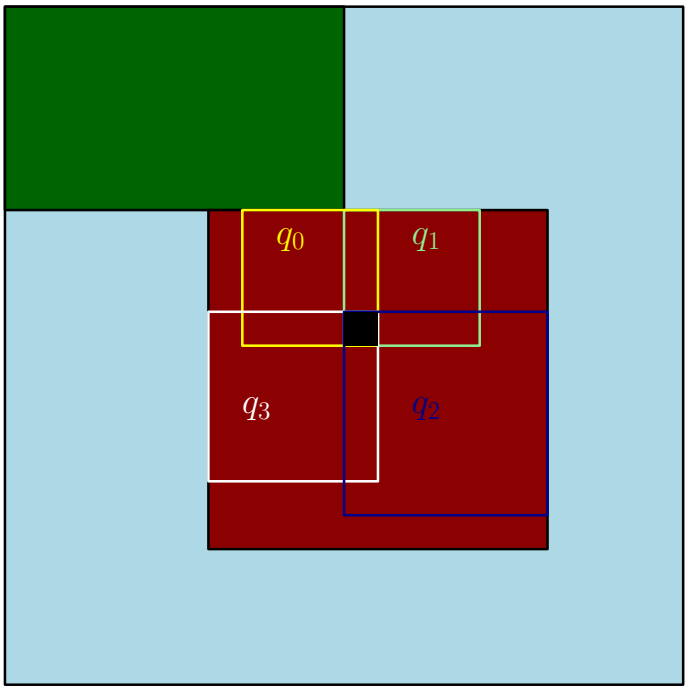

FIGURE 4.3: An illustration to show how the quadrants can grow around the test pixel. The green, red and blue shaded areas simulate different colour patches in an image. Quadrants $q_{0:3}$ are the Adaptive Kuwahara quadrants. The blue quadrant q_2 grows to a final size of 6×6, whereas q_3 grows to a final size of 5×5 and quadrants q_0 and q_1 grows up to 4×4.

smaller, the process is repeated until the variance is larger than the previous variance. When all areas stop increasing in size, the mean of the area with the smallest variance is used as the new data point. Bartyzel states that the new algorithm is much faster than the adaptive median filter while removing the need to predetermine a window size just as the original Kuwahara filter [54].

The line noise present in our data sets would greatly affect the result of our segmentation stage because of the significant difference in the spectral values they contain. The first criteria is to remove the lines as best as possible. A 3×3 Median Filter suffices at removing these lines. The reason why a median filter was chosen is because this noise would present high values at the extremes of a histogram. The second criteria is to smoothen the noise generated from the random effect of the camera sensors. The Adaptive Kuwahara Filter is used to smoothen the random noise while preserving the important edges.

After applying these steps, it is important to show that the noise has been greatly reduced, and information on the edges has not been lost. To simulate the effect of our filtering stage, an RGB form of our data is presented. A singular line of data, as shown in Figure 4.4(a) is used to extract intensity values; if the pigment is green we use the green

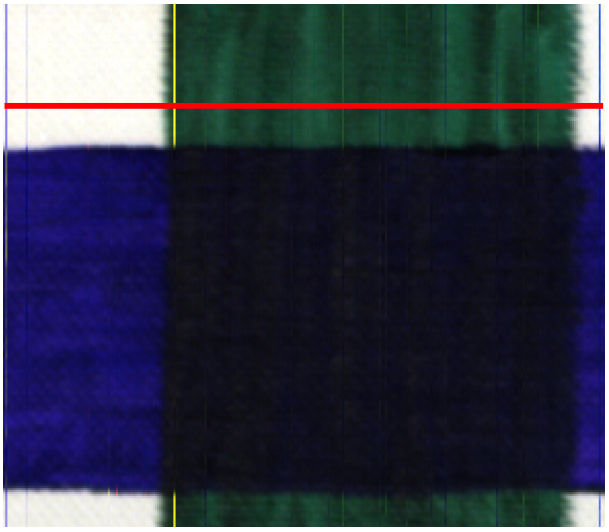

(a)

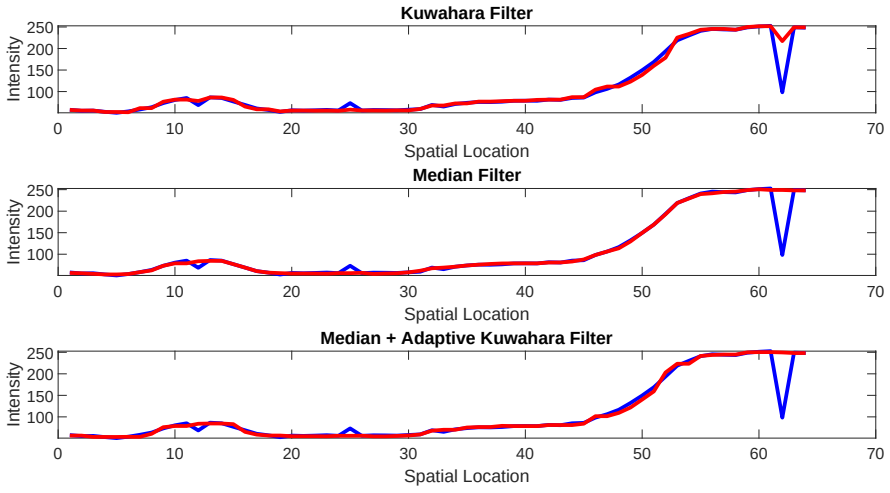

(b)

FIGURE 4.4: Plots in (b) compare the Median and Adaptive Kuwahara filters on sample image in (a). The red signatures represent the red strip of pixels from the filtered image, whereas the blue signatures represent the red strip of pixels from the unfiltered image.

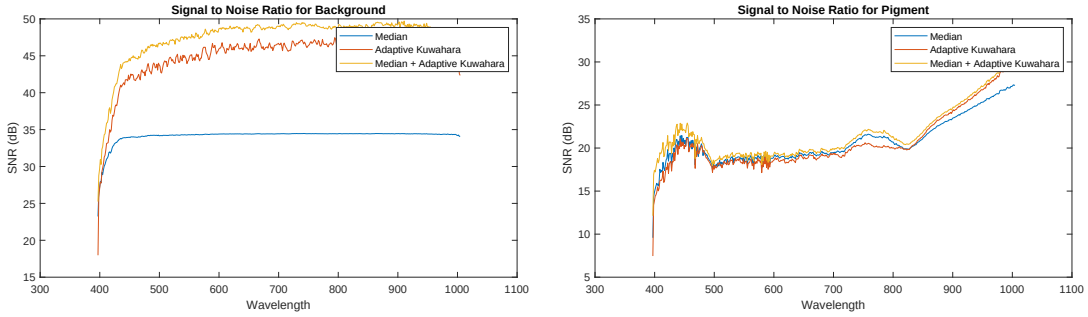

(a) (b)

FIGURE 4.5: SNR at each wavelength of the datacube for two ROIs. (a) A canvas (substrate) region, (b) A pigment region.

channel intensity values. The resulting outcome is shown in Figure 4.4(b). Additionally, SNR values for each filter on each wavelength image are calculated to make sure that the RGB simulation is effective. Since we do have ground truth images, the reference image can be assumed to be the filtered image. Then the SNR will be the signal power divided by the noise power. However, from Figure 4.5 one can see that the stand-alone median filter does not perform well, but it aids the Adaptive Kuwahara filter to improve the SNR by removing the line noise first. Although the camera itself is stated to have an SNR of 600:1, if we zoom in the image we can still notice some level of noise between neighbouring pixels. Hence the importance of an arithmetic mean filter, in our case the Adaptive Kuwahara filter, which preserves edges.

4.3.2 Filtering Pixel Vectors

Remembering that our data is 3-D in nature, spectral filtering would also help reduce the effect of random noise between neighbouring wavelength images. This type of noise is caused by the quantum differences in each of the sensors in the camera. It causes unwanted oscillations over the spectral signature. It is important to remove this noise to enhance prominent features between different signatures. To remove these slight artifacts caused by differences in sensor sensitivity, a centered moving average filter is applied. The centered moving average filter is used with a window size of nine. Other filters such as Savitzky-Golay filters exist, however, the moving average is sufficient enough for our application [27]. Additionally, the first 20 initial and final 20 channels appear to be noisier than others and are therefore truncated.

By filtering the raw data, the image cube can be processed further with less interference from the previously present noise. In the next section, we will address a feature extraction method that segments regions by the similarity in their spectra.

4.4 Feature Extraction

In HSI, feature extraction can be performed by looking at each and every individual pixel. This method appears to be a very common approach as shown in many other research [1, 5, 28]. Although it accurately identifies each pixel, the method creates noisy results because of outliers and boundary artifacts. One way for reducing this problem is by including a segmentation step between the filtering and feature extraction stages. Segmentation does not only reduce the chance of detecting outliers by grouping them with distinct features, but can also provide a visual aid to the analyst as to where the borders of adjacent pigments are.

Although the segmentation field has been widely explored in many different domains as seen in Chapter 3, the Simple Linear Iterative Clustering (SLIC) proposed in [45] provides good boundary adherence while retaining compact superpixels. Moreover, this superpixel technique can be adapted to hyperspectral images as seen in the implementation proposed in [4]. However, we notice that although the segmentation manages to capture those regions of similar features, it also suffers greatly from over-segmentation. Moreover, one must be careful not to under-sample, as small regions with distinct features may not be captured. Therefore, the authors had to compromise between over-sampling and capturing small features. In this section, a method to overcome the problem of over-sampling while managing to capture small features is proposed.

4.4.1 Paint Segmentation

In this section, we propose a novel region-merging based algorithm for hyperspectral applications called Spectral Similarity Merging (SSM). The aim is to separate different regions by observing their spectral similarity, that is the similarity between the spectral signatures of neighbouring regions. Moreover, it is very important to keep a smooth transition in regions between adjacent pigments while keeping a low over-segmentation. The problem is split into three main steps:

- Region Merging of similar neighbours.
- Pixel re-assignment of regions having a high intra-class variance, and Perimeter Smoothing of low intra-class variance regions.

- A final Region Merging checks for nested regions or newly created similar regions.

In the Region Merging step, neighbours having similar spectra are merged together, updating their spectra. Those regions that end up having dissimilar spectra remain unmerged. In step 2, these regions are checked for high intra-class variance. A high intra-class variance means that there are multiple spectral properties within the region, which suggests that the region must be split. Therefore, we re-assign the pixels to low intra-class variance neighbouring regions having similar spectra. This step follows Ward's Criterion where the variance of a region is minimised by reassigning pixels to other regions with minimal effect on the variance of the candidate region. Once these regions have been processed, we analyse the low intra-class variance regions by dragging a window over the perimeter of the regions and analysing the pixels that fall under this window. These aforementioned steps may cause regions to have an improved homogeneity, and hence the possibility for them to be spectrally similar to their neighbour, therefore, we perform a final merging check.

The initialisation of the algorithm is inspired from [45] in which a regular square grid is superimposed over the image, where each square is a superpixel. Achanta et al. suggests that the size of each square in the square grid is calculated using $S = \sqrt{\frac{N}{k}}$, where S is the dimension of each square, N is the number of pixels in the image and k is the number of superpixels. However, this creates a problem in which the analyst must qualitatively inspect different results for a different number of superpixels, and decide which value returns the best segmentation. Instead, we can inspect the image prior to choosing k and look for the smallest region. If this region is important and must be captured, then according to Nyquist's criterion we must sample the region at, at least twice this smallest dimension. We achieve this by setting S to be equal to twice the width of the bounding box enclosing the region. By changing the subject of the formula we obtain the number of superpixels as $k = \frac{N}{(2S)^2}$ and by doing so, we eliminate any unnecessary over-segmentation, while minimising under-segmentation. In Figure 4.6, a block diagram illustrates the entire process of our algorithm.

4.4.1.1 Region Merging

Every square in the initial grid described above is assigned a spectral centroid by taking the average of the entire spectrum. The spectral centroid is used to obtain a similarity measure between two neighbouring regions. We use one of the features of Region Adjacency Graphs to obtain a list of neighbours for each region. Before we go any further, we must understand what kind of spectra we consider to be similar. There are

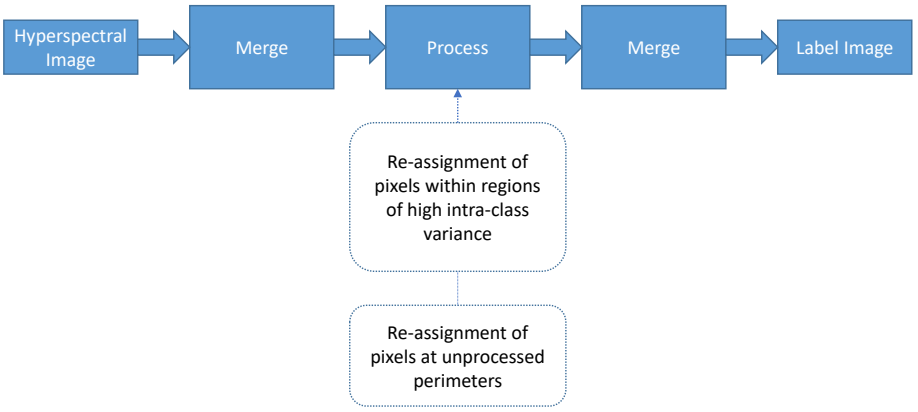

FIGURE 4.6: A block diagram of the segmentation process for a raw input to obtain the output. The white boxes briefly explain the process within the main blocks.

three types of similarity metrics: intensity similarity, shape similarity, and intensity and shape similarity. Intensity similarity metrics are those like Euclidean distance, where a distance is measured based on the squared difference of two values for the same point. Shape similarity metrics like Spectral Angle Mapper (SAM) [55] and Spectral Correlation Mapper (SCM) [49] are based on spectral angle and spectral correlation, respectively. Whereas others such as Spectral Similarity Scale consider both similarities in shape and intensity [56]. Since we define spectral similarity to be based only on the shape for this application, we choose to continue on the basis of correlation. What this means is that, if a region containing the same pigment has varying intensities due to differences in brush pressure during application, the difference will be ignored. The reason why SCM is chosen over SAM is because although they are both spectral signature feature measures, the former is not affected by spectral gain nor offset, whereas the latter is affected by spectral offset [56]. The Kernel-SAM [47] is a modified version of SAM in which the overall accuracy, when used for classification purposes, increased by 3.5%. The equation for this metric can be found in Chapter 3. Since this extended metric can improve the accuracy of its predecessor by 3%, we apply the theory to the SCM and obtain a Kernel-SCM, which aims to improve the accuracy for our application. The equation for this metric is given by Equation 4.1:

$$\alpha = \frac{\sum_{i=1}^{\lambda} \mathbf{\Phi}(t_i - \bar{t}) \cdot \mathbf{\Phi}(r_i - \bar{r})}{\sqrt{\sum_{i=1}^{\lambda} (\mathbf{\Phi}(t_i - \bar{t}))^2} \sqrt{\sum_{i=1}^{\lambda} (\mathbf{\Phi}(r_i - \bar{r}))^2}} \tag{4.1}$$

where Φ is the kernel function, t is the test reflectance, r is the reference reflectance,

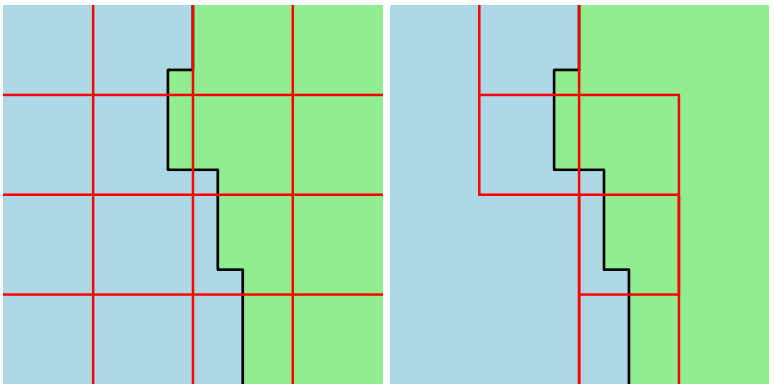

(a) (b)

FIGURE 4.7: Demonstration to show how a superpixel having high similarity merges with another superpixel. Light blue superpixels are merged together and light green superpixels are merged together. Superpixels having different features should not merge.

$\bar{t}$ and $\bar{r}$ are the means, λ is the number of wavelength channels and α is the similarity coefficient.

The merging criterion is solely based on spectral similarity, meaning no spatial features are being considered at this stage. This is important because we do not want to limit the ability of growth of a region, but we do want to restrict the growing to remain homogeneous in terms of spectra. Two regions are allowed to merge if the spectral similarity is lower than a given threshold, and once a region merges it cannot merge a second time during the same iteration. This is done to prevent a region from growing out of proportion, which may cause a large change in features in comparison to smaller regions. The outcome of this stage should minimise the total number of regions by merging those regions that have homogeneous properties, preparing the data for the next stage. A simple example is illustrated in Figure 4.7

4.4.1.2 Pixel Re-assignment

Regions that fail to satisfy the merging criterion remain unmerged. These regions may fail to satisfy the merging criterion for two reasons. The first reason is because their similarity was borderline with the threshold, and will be explained further in Chapter 5. The second reason is because the region contains multiple spectral features. The multiple spectral features cause large changes in the spectral signatures of the centroids, and therefore, an increase in the intra-class variance of these regions. Since we are dealing with spectral features, and an image is available for each wavelength, the spectral variance must be a vector. The variance vector captures spectral differences and allows us to identify high intra-class variance regions by analysing the variance signature. In Figure 4.8, an

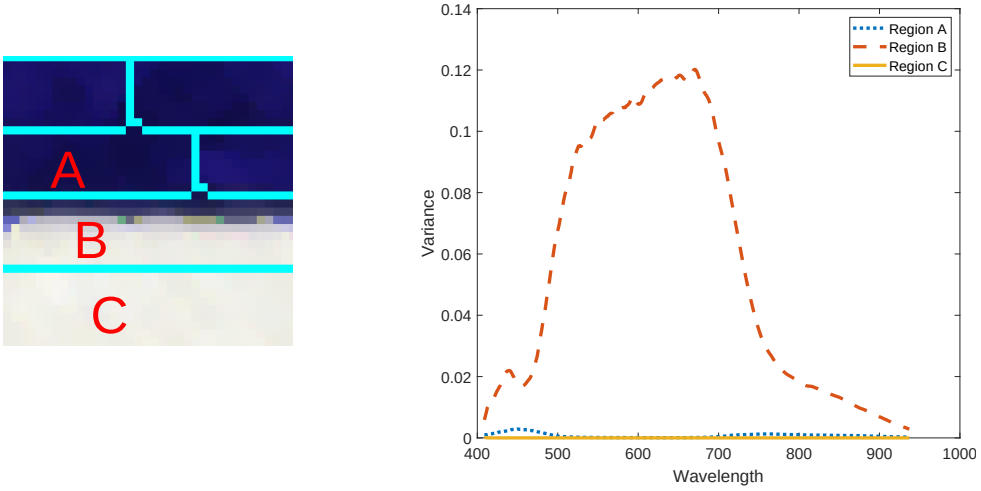

(a) (b)

FIGURE 4.8: Showing the difference in intra-class variances for a homogeneous blue region A, a mixed Region B and a homogeneous white Region C. The graph shows a low variance throughout the entire spectrum in both homogeneous regions.

illustration shows the difference in spectral variances between homogeneous regions and mixed regions.

Once a high-intra class variance region has been identified, the region goes through a pixel re-assignment process where each pixel is re-assigned to a neighbouring region that is considered to have a low intra-class variance, as illustrated in Figure 4.9. This stage considers both spectral and spatial features. The spectral features are important such that a pixel is not re-assigned to a spectrally different region, whereas the spatial features allows us to have some control over the pixel re-assignment. Although the spectral features are compared using the centroid, the spatial features are compared using the single-link method. The spatial features are obtained by computing Equation 4.2:

$$v_i = \min \left(\sqrt{\sum_{j=1}^{n} (\boldsymbol{p} - \boldsymbol{q}_j)^2} \right) \tag{4.2}$$

where n is the total number of pixels for region i, $\boldsymbol{p}$ and $\boldsymbol{q}$ are the vectors containing the x and y coordinates of the pixel under re-assignment and the pixels from candidate regions, respectively.

The advantage of using the single-link for this application is because the low intra-class regions will have irregular sizes following the Region Merging stage, and therefore if we use the complete-link a bias toward the smallest region is created. To obtain a

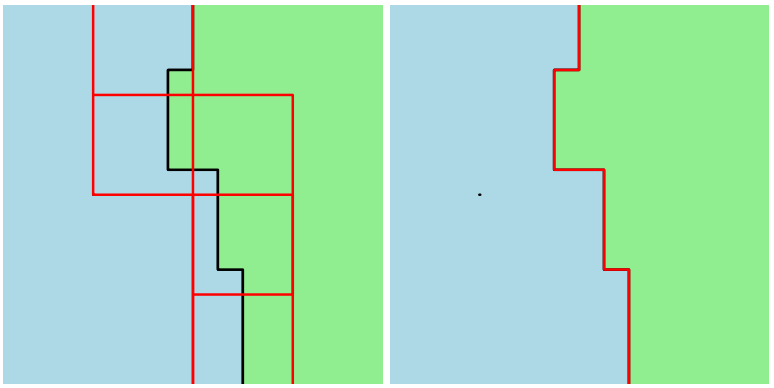

(a) (b)

FIGURE 4.9: Demonstration to show how superpixels that have high intra-class variance have their pixels re-assigned to neighbouring superpixels.

single distance metric, we combine both the spectral similarity and the spatial distance terms into one metric as done in SLIC [45]. The distance metric is given by Equation 4.3:

$$D_i = \sqrt{\frac{v_i}{\sum_{j=1}^{n} v_j} + \alpha\left(\boldsymbol{p}, \boldsymbol{c}_i\right)} \tag{4.3}$$

where α is the Kernel-SCM, $\boldsymbol{v}$ is the spatial distance vector, $\boldsymbol{p}$ is the spectral centroid of the pixel, $\boldsymbol{c}$ is the spectral centroid of the candidate neighbouring region, i is the index for the available candidate regions and n is the total number of candidate regions. The spatial vector $\boldsymbol{v}$ contains the minimum distance between the pixel being analysed and all pixels within a candidate neighbour, and as can be seen in the Equation 4.3, it is being normalised.

Once all regions having a high intra-class variance have been processed, we must process the low intra-class variance regions. Note that due to the previous stage, high intra-class variance regions should now be transformed to low-intra class variance regions. The processing here is less complicated because we already know that the region is homogeneous, but we are unsure if the perimeter is accurate since the region has a regular shape. A window of size $\frac{S}{2}$ is passed over the perimeter of these regions, and each pixel that falls under the window is checked for spectral similarity against all different spectral centroids for all labels that are present under this window. The re-assigning criterion used here is equivalent to that used in Stage 1. The result of this stage should present a smooth region segmentation with no fragmentation at adjacent pigments.

4.4.1.3 Checking for Nested Regions

In the final stage, the resulting regions are checked for similarity using the same merging criterion in Stage 1. The reason why we should perfom a secondary region merging step is because in Stage 2, the regions have been processed and their spectral centroids are updated using only homogeneous spectra. Therefore, any regions that had borderline similarity in Stage 1 may now fall within the merging criterion and can be merged. Additionally, there may also be regions that appear as nested regions which may also be created by the processing steps in Stage 2. Although not always the case, these regions may contain the same spectral features as their neighbour, hence, we can evaluate the similarity at a slightly more relaxed threshold.

Once this process has been completed, a post-processing step cleans the label image from any spurs and performs a closing operation to remove regions made of a single pixel because we assume that it is not practical that a single pixel-wide dissimilar feature exist. In Figure 4.10 an illustration of the entire process for a typical sample is given, in which one can observe that the regions respect the boundary adherence problem, while minimising the over-segmentation problem. Further analysis regarding the results for such metrics can be found in Chapter 5.

4.4.1.4 Thresholds

This algorithm is based on finding the optimal threshold for the region merging steps. Finding the optimal threshold is important because we want to achieve the best rate between true positives and false positives. Since some pigments are very similar to each other, for example those illustrated in Figure 4.11, we must base our threshold on separating these two paints. We note that in order to be able to separate these two pigments, we may have to accept some trade-off when pigments are dissimilar. This will cause some degree of over-segmentation, but still less when compared with other state-of-the-art methods as discussed in Chapter 5.

To find the optimal threshold, we build a receiver-operating characteristic (ROC) for a set of thresholds ranging between 0.05 and 0.9 in steps of 0.05, measuring the true positive and false positive rates for each step. This task is carried out for two samples: a sample containing dissimilar paints, and a sample containing the two similar red paints, as seen in Figure 4.12. In sample 1, where the pigments used are very different, the AUC (0.88) is higher because the threshold manages to separate the two pigments even at low values. On the other hand, for sample 2, we observe that ROC curve the pigments can be separated at the cost of lower AUC (0.84). Therefore, from these plots we can

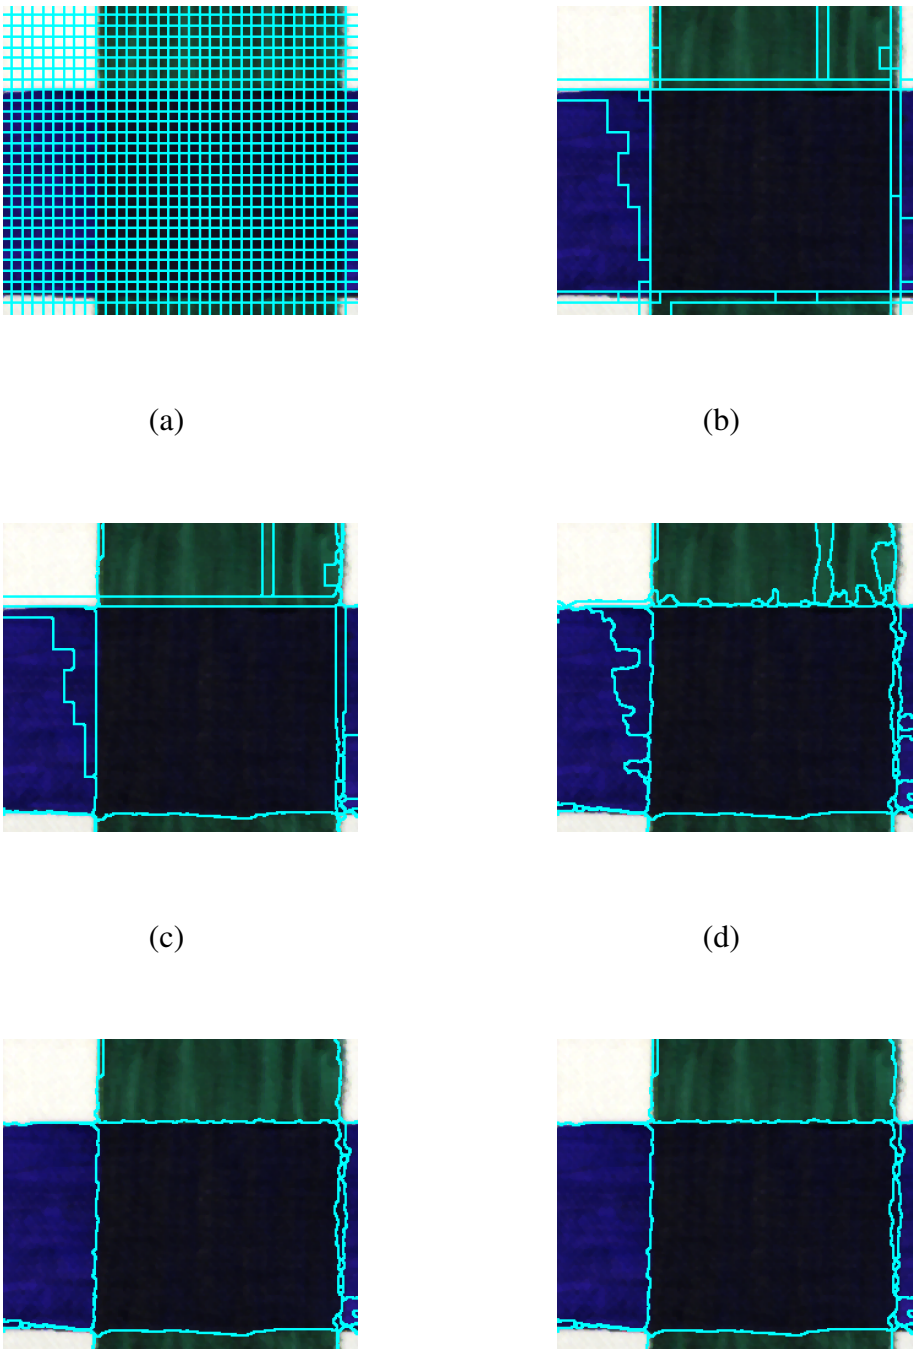

(a) (b)

(c) (d)

(e) (f)

FIGURE 4.10: Process as seen by showing the results from each stage. (a) Initialisation stage, (b) Region Merging stage, (c) High intra-class Variance Pixel Re-assignment stage, (d) Low intra-class variance Perimeter Smoothing stage, (e) Checking for Nested Regions stage, (f) Post-Processing stage.

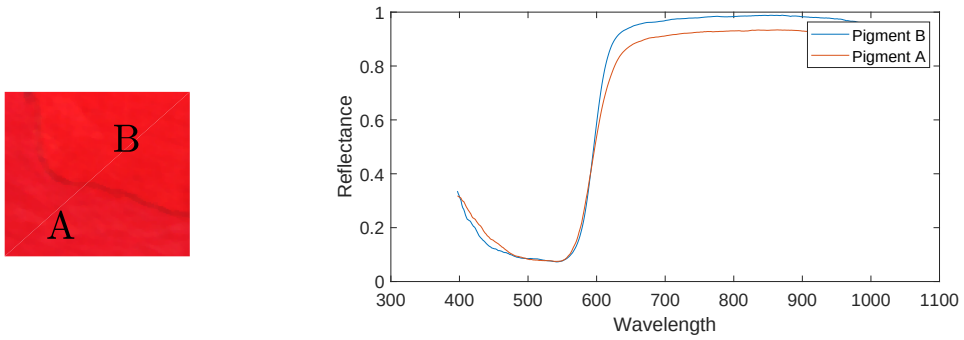

(a) (b)

FIGURE 4.11: Illustrating the similarity between the two red pigments Permanent Geranium Lake Red (A) and Cadmium Red (B). (a) Is an RGB representation of the two pigments taken from one of our samples, (b) illustrates the reflectance curves of the two pigments.

qualitatively conclude that the best threshold value that separates similar paints is 0.45. The same procedure has been carried out to extract the best threshold for the Pixel Re-assignment stage.

In this section we proposed SSM, a region-based segmentation algorithm. As we have already seen the algorithm is initialised by creating a regular grid in which each square is considered to be a region having a spectral centroid. Neighbouring regions are then merged depending on their spectral similarity. Regions that contain two different paints will remain unmerged, and a Pixel Re-assignment stage further splits these regions by merging their pixels with low intra-class variance neighbours. When no high intra-class variance regions remain, the perimeter of the low intra-class variance regions is processed. This leads us to perform a final Region Merging stage followed by a post-processing step that performs closing and spur removal. The algorithm provides a strong boundary adherence and a low over-segmentation solution. In the following section, three methods are proposed to analyse paint mixtures in the VNIR region.

4.5 Paint Classification

In the previous section, the different features of a painting have been extracted in the form of homogeneous regions or superpixels. After segmentation, a region will consist of homogeneous spectral signatures. The goal is to identify these spectral signatures

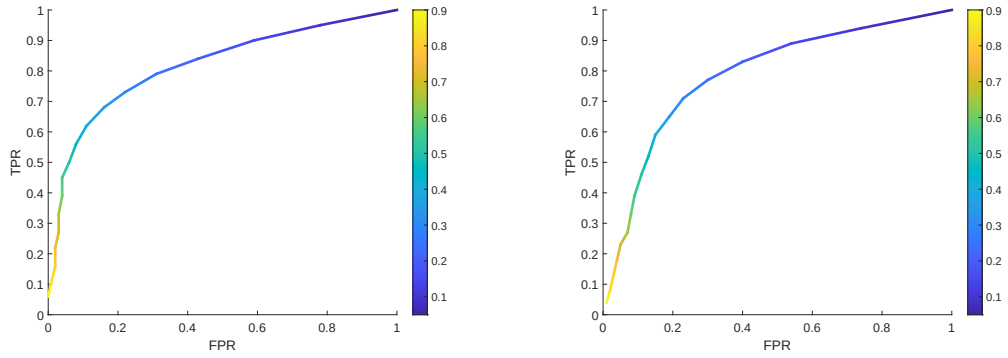

(a) (b)

FIGURE 4.12: Two ROCs for the two samples used to determine the right threshold. (a) Sample 1 in which the pigments are dissimilar in shape having an AUC = 0.88 (b) Sample 2 in which the pigments are similar in shape having an AUC = 0.84. The colourbars represent the threshold values.

based on their spectral signature in order to label the region based on the paint that it contains. In our work, a homogeneous region can be based either on a single paint, on a mixture of a paints, or on paints layered over each other. Paint mixtures can be made from multiple pure paints, typically not more than three paints. For example, an artist may decide to create a green shade by mixing a blue and a yellow, and then they can add a black or a white to darken or brighten the shade. In this dissertation, we will focus on paints made from two pure paints, and we will not consider the addition of black or white colours to these mixtures.

Our classification problem can, therefore, be split into two: pure paint regions, and mixed paint regions (including layered paints). To separate regions made from pure paints from regions having a mixture of paints, we can classify all regions directly, and identify those regions in which the similarity is too small because of the dissimilarity between the features of pure paint references and paint mixtures, marking the region to contain a mixed paint. Consequently, we can then focus on mixed paints after the pure paint regions have been classified. The objectives of this section can then be formulated as:

- Given an unknown colour signature made from one paint, the correct name of that pigment must be obtained.
- Given an unknown colour signature made from more than one paint, the correct class of the paints making the mixture must be identified, preferably also identifying the specific paints.

- Alternatively, if a mixed paint is defined as an individual paint, the correct paint mixture must be identified with irrelevance to mixture variants.

In this section, we design three different methods i) to identify the pure and mixed paints, ii) to compare the results obtained from a local and global NNMF optimisation problem, and iii) to analyse the effects of signal variance in the global optimisation problem. The three methods will include the use of NNMF. In the first method, Hierarchical Paint Analysis, a typical similarity classifier will be used to determine the pure paint regions, while NNMF will be used on the unclassified regions to solve a local optimisation problem for a select combination of paints. The method follows a hierarchical approach in which first, the correct classes of paints are determined, then the specific paints are determined. The four different classes of paints are: Yellow, Green, Blue and Red. In the second method, Global NNMF for Paint Analysis, NNMF will be used to solve a global optimisation problem in which no prior knowledge regarding the paint in the region is known. In the third method, Direct Classification of Mixtures, the mixed paints will be treated as individual paints, and the goal here is to see whether classification of mixtures improves when reference spectral mixtures are used.

4.5.1 NNMF for Paint Analysis

For all methods mentioned, which will be further discussed in the following subsections, we use the same notations for NNMF. In particular, $\boldsymbol{I}$ is the data input matrix, $\boldsymbol{H}$ is the feature matrix, and $\boldsymbol{W}$ is the weight matrix. When NNMF is applied to $\boldsymbol{I}$, the basis matrix $\boldsymbol{H}$ will contain the prominent features for each variable in the input data, based on the selected rank k, whereas the weight matrix $\boldsymbol{W}$ will contain the weight coefficients for each data point in the input data, also based on the selected rank. From literature, we know that the mixing of paints can be assumed to be linear, in fact, many papers use the linear unmixing model $\boldsymbol{I} = \boldsymbol{H}\boldsymbol{w} + \boldsymbol{e}$ and optimise the cost function in Equation 4.4 [5, 18, 51].

$$E = \|\boldsymbol{I} - \boldsymbol{H}\boldsymbol{W}\|^2 \tag{4.4}$$

In our work, we use NNMF to optimise for $\boldsymbol{W}$ with a known $\boldsymbol{H}$, by optimising Equation 4.4 using the multiplicative update rules defined by Equation 3.11 and Equation 3.12 in Section 3.2.3.2. Since the general form of NNMF does not keep a uniform scale during factorisation, we can make use of a sum-to-one constraint on $\boldsymbol{W}$ such that the weight coefficients for each feature in a data point sums to one. Since the multiplicative update rules make use of the input data to optimise $\boldsymbol{H}$ and $\boldsymbol{W}$, we cannot make use

of a sum-to-one constraint within the update rules to restrict the values in $\boldsymbol{W}$. Instead, we can make changes to $\boldsymbol{H}$ and $\boldsymbol{I}$ to constrain the values in matrix $\boldsymbol{W}$ to sum-to-one for each data point, as given by Equation 4.5 and Equation 4.6, respectively [57]. The sum-to-one constraint will allow us to compare the weights with the actual ground truth mass-to-mass ratios more easily.

$$\mathbf{H} = [H_{r,1}, H_{r,2}, \ldots, H_{r,2\lambda}, \delta 1_{r,1}]' \tag{4.5}$$

and,

$$\mathbf{I} = [I_{N,1}, I_{N,2}, \ldots, I_{N,2\lambda}, \delta 1_{N,1}]' \tag{4.6}$$

where δ is a manually set parameter to control the constraint (parameter varies according to data), λ is the total number of wavelength indexes, r is the reflectance value, and N is the data point or pixel. The result from this matrix factorisation for $\boldsymbol{W}$ should give us an indication of how much of each feature per wavelength in $\boldsymbol{H}$ is required to create $\boldsymbol{I}$. Therefore, from the reconstruction, we can identify the best two classes that re-create the mixture in $\boldsymbol{I}$.

We also know that assuming the reflectance of the mixture to be linear can be highly inaccurate because many factors are ignored, as mentioned in Chapter 2. In [20], the authors applied the KM Theory to learn the absorption and scattering coefficients for the reflectance data. We follow the approach in [20] and observe the difference between using reflectance, absorption and scattering coefficients.

4.5.2 Normalisation

In the literature, we found that for better accuracy in determining mixture ratios one can normalise the data using continuum removal [58, 59]. The reason why continuum removal improves the accuracy is because it enhances features, such as absorbance bands. Hence, subtle changes in absorbance features are more easily detected. Therefore, if a signal contains two separate absorption bands representing two pure paints, but the layer thickness varies - causing a difference in amplitude - continuum removal will normalise this offset, allowing a better analysis of the absorption bands, especially the weaker absorbance bands [59]. In our work, we find the continuum of a signal by joining lines from local maxima to the next. By solving Equation 4.7 [58] we can calculate the continuum removed signal to obtain a normalised spectrum:

$$CR = \frac{S}{C} \tag{4.7}$$

where S is the signal, C is the set of local maxima representing the continuum, and CR is the resulting continuum removed signal.

From preliminary experiments, we find that applying the continuum removal technique improved correlation results by 1% and NMF results by 11%. The data used for these experiments was taken from our own sample images. Following these observations, we will not be including any results for non-normalised data in Chapter 5.

4.5.3 Reflectance to KS Coefficients

As discussed in Chapter 2, paint mixtures can occur either by mixing paints or by layering one paint over the other. Layering paints over each other is more complex to analyse, because the thickness and positioning contribute toward the reflective measurement [17]. Mixing paints, on the other hand, will result in two paints being in the same layer and hence only the layer thickness and mixing ratios contribute toward reflective measurements. A paint mixture is normally given by a linear model as shown by Equation 4.8:

$$\left(\frac{A}{S}\right)_m = \frac{\left(1-\sum_{i=1}^{n} c_i\right) a_t + c_1 a_1 + c_2 a_2 ... + c_n a_n}{\left(1-\sum_{i=1}^{n} c_i\right) s_t + c_1 s_1 + c_2 s_2 ... + c_n s_n} \tag{4.8}$$

where A/S is the absorption and scattering ratio, m represents the mixture of the n pigments, c denotes concentrations by weight of pigments 1 and 2, a and s are the absorption and scattering coefficients respectively, and t is the white tint. In cases where a white tint is not present in the mixture, the formula is reduced to Equation 4.9 [17].

$$\left(\frac{A}{S}\right)_m = \frac{c_1 a_1 + c_2 a_2 ... + c_n a_n}{c_1 s_1 + c_2 s_2 ... + c_n s_n} \tag{4.9}$$

Finding a library for absorption and scattering coefficients that included all of the pigments that we are using, at exactly the right wavelengths, and with the same sample properties was difficult. Therefore, we chose to learn these coefficients by finding a best possible starting point from an available dataset and match the wavelengths by reducing the sample size of our data. The Okumura dataset [19] includes 26 pigments and their corresponding reflectance, and scattering and absorption coefficients in the range between 400nm-750nm in steps of 10nm (36-channel). Since the database is being used only as a reference, or starting point, we need to use a non-linear least squares

method to find the best fit for an absorption and scattering coefficient that best represents the input reflectance [20], given by Equation 4.10:

$$R = \frac{1 - \xi\left(x - y \cdot \coth\left(yst\right)\right)}{x - \xi + y \cdot \coth\left(yst\right)} \tag{4.10}$$

where,

$$x = 1 + \frac{a}{s}$$
$$y = \sqrt{x^2 - 1}$$

where t is the thickness, ξ is the substrate's signature, and R is the total reflectance. For turbid media (opaque), the thickness can be assumed to be infinite, which simplifies Equation 4.10 to Equation 4.11 [19].

$$R_\infty = 1 + \left(\frac{A}{S}\right) - \sqrt{\left(\left(\frac{A}{S}\right)^2 + 2\left(\frac{A}{S}\right)\right)} \tag{4.11}$$

Therefore, we can use Okumura's dataset [19] to learn scattering and absorption coefficients for our own ground truth reflectance spectra.

In Figure 4.13, the scattering and absorption coefficients for each ground truth are illustrated. After learning the absorption and scattering coefficients, the four classes of paints are created by taking the average of each paint in that class. Moreover, after learning, we should be able to apply the coefficients to the KM equations and obtain R, to make sure that there are no errors in learning. In Figure 4.14, we show an example of two pure pigments and their corresponding mixture. By adding the numerator and denominator of Equation 4.9 separately, we obtain a single signature for both absorption and scattering representing a mixture. We can then use these mixture coefficients to obtain the mixture's reflectance using Equation 4.10 and Equation 4.11. On inspecting Figure 4.14, we notice that there is not much difference between the two equations mentioned in Figure 4.14 except for a slight difference in spectral intensity.

When comparing identification results for the absorbance and scattering coefficients, against identification results for reflectance, it was observed that the absorbance and scattering coefficients also improve correlation and NMF results. It has been observed through these initial experiments that correlation results improves by 1% and NMF results improve by 2%. Therefore, similarly to the decision made for the normalisation procedure, we will not be including reflectance results in Chapter 5. It is also stated in [20] that the absorbance and scattering coefficients should improve the separability between paints when compared against the same signal in the form of reflectance.

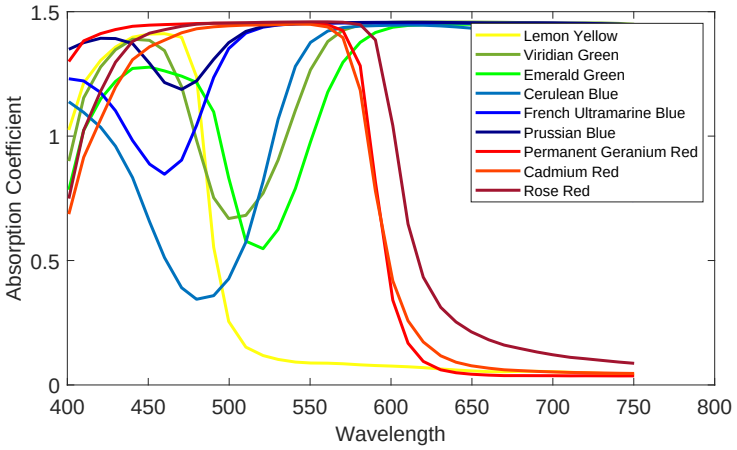

(a)

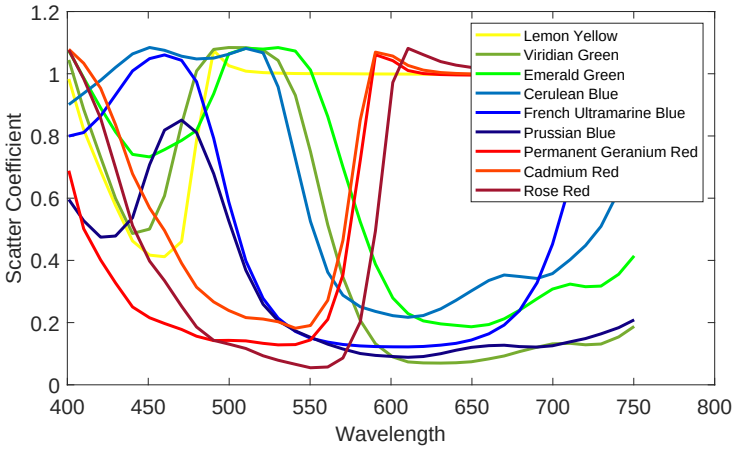

(b)

FIGURE 4.13: Illustrating the learned (a) absorption (b) scattering coefficients for the nine paints in our dataset.

4.5.4 Method 1 - Hierarchical Paint Analysis

One of the main hurdles in paint identification is not the identification of individual paints, but rather the identification of regions that may be made up of a mixture of individual paints. Therefore, one of our goals, is to design a method that analyses these

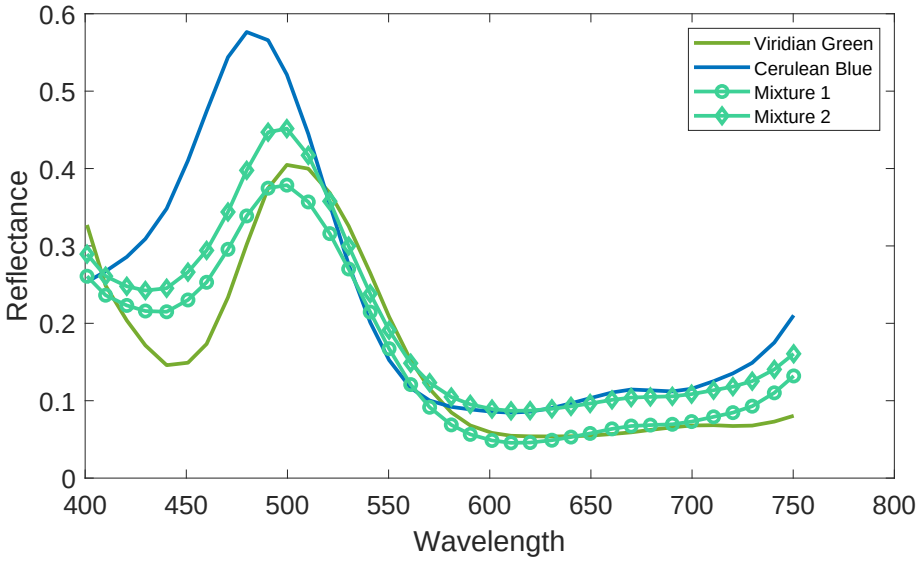

FIGURE 4.14: A mixture for Cerulean Blue and Viridian Green is simulated using Equation 4.11 (Mixture 1) and Equation 4.10 (Mixture 2). Mixture 2 assumes a thickness of 1, whereas mixture 1 assumes infinite thickness and hence the reduction in amplitude.

regions on a painting, and determines any regions that may be mixtures. For this to be possible, the palette containing all the individual paints used for the painting must be available. In this section, we use a modified version of the Kernel Spectral Angle Mapper (KSAM) algorithm which we introduced in Section 4.4.1.1 to identify pure paint regions. For mixed regions, we apply the well known NNMF [51] algorithm to the problem of pigment unmixing. Although the algorithm has been used in previous works to identify concentrations of known pure elements in a mixture [18, 60], we attempt to apply the algorithm to a blind problem. This means that, prior knowledge regarding the mixture is not used, but rather aim to identify what is in the mixture based on the knowledge of the palette of individual paints used on the painting.

The different paints in our library are introduced in Section 2.2. The signatures of these paints are shown in Figure 4.15. In this work, the mixed regions in the mock-ups are made up of two paints. Therefore, our method is based on the prior knowledge of knowing that the maximum number of paints in the mixture is two.

We split the problem of pigment identification into two stages. Initially, all regions are checked for high similarity with pure paints in the library. The second stage only

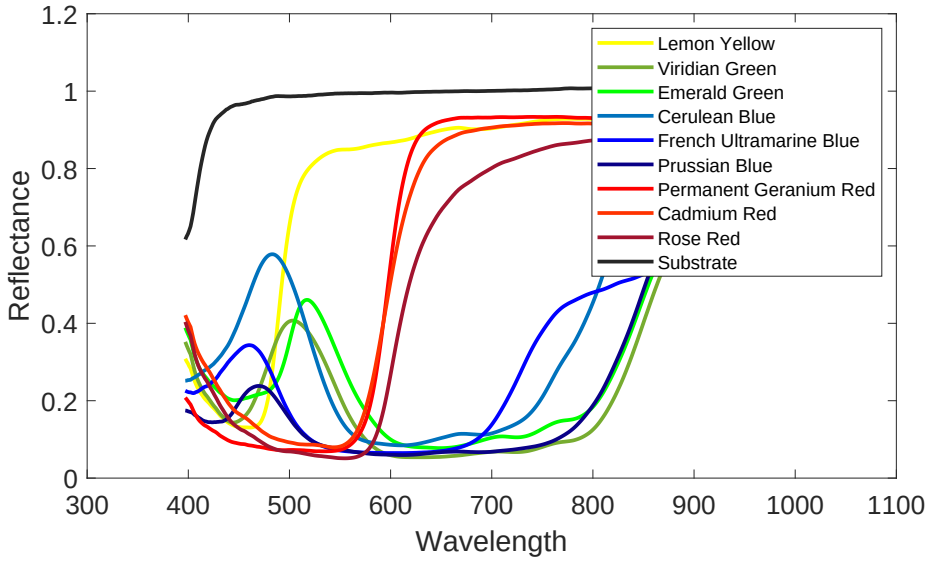

FIGURE 4.15: An illustration for the paint reflectance spectra present in the spectral library. The substrate marked in black is also shown.

applies to those regions which do not present a high similarity to any paints in the library. In this second stage the pure paints are grouped according to their paint classes, that is, all red paints are grouped into the red paint class and likewise for the other paints resulting in four paint classes.

Regions that are not classified as pure regions pass through a hierarchical two-stage pipeline, which first checks the signature against the paint classes. A second stage then computes the best similarity between the mixture and the individual paints in each identified paint class. In the following subsections, an in-depth explanation of the method is given. A block schematic in Figure 4.16 explains this process in more detail.

4.5.4.1 KSCM Paint Identification

Paints are made up of different minerals and elements, each of which is represented by a specific signature in the Visible-Near Infrared (VNIR). A closer look at the different classes of pigments by observing Figure 4.17, shows that the red class has much less intra-class spectral variance when compared to blue and green. This means that it will be much harder to separate red paints in the VNIR range [1]. Moreover, one must

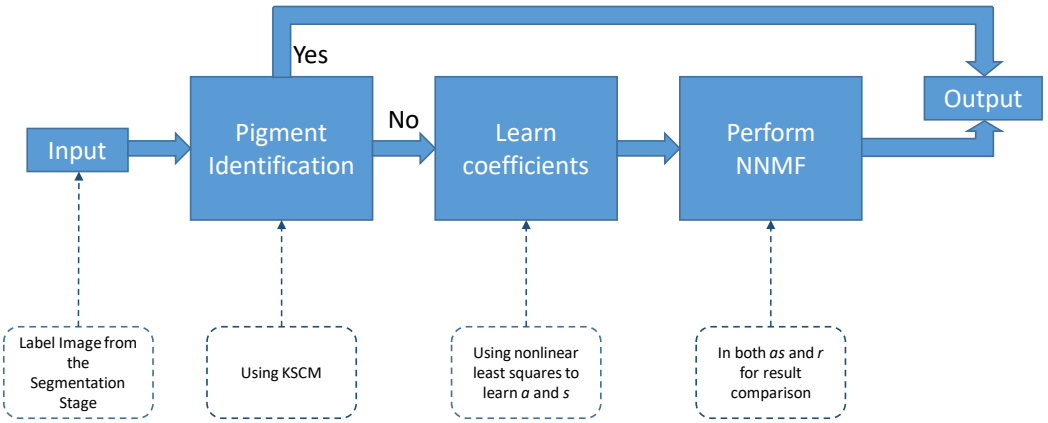

FIGURE 4.16: A block diagram that illustrates the pipeline for the classification process of a segmented image. The white boxes briefly explain the process within the main blocks.

understand that in paint that is not fully turbid, such as oil-based paint and 'watercolours', spectral variance plays a huge role especially at areas in which the paint layer thickness varies because more light reaches the substrate. Consequently, techniques such as SAM or KSAM [47] cannot be used as a change in intensity will affect the similarity result, causing inaccuracies throughout all classes. On the other hand, although SCM is better since it removes the mean, it also results in small inaccuracies when dealing with the red pigments. Therefore, we apply the KSAM approach to SCM, given by Equation 4.12. In [47] the authors mention that the kernel can be polynomial, sigmoidal or a radial-basis function. From our experiments we notice that a radial-basis function provides the best results for separating pigments, including red pigments.

$$\alpha(t,r) = \frac{\sum_{i=1}^{\lambda} \mathbf{\Phi}(t_i - \bar{t}) \cdot \mathbf{\Phi}(r_i - \bar{r})}{\sqrt{\sum_{i=1}^{\lambda} (\mathbf{\Phi}(t_i - \bar{t}))^2} \sqrt{\sum_{i=1}^{\lambda} (\mathbf{\Phi}(r_i - \bar{r}))^2}} \tag{4.12}$$

where Φ is the kernel function, t is the test reflectance, r is the reference reflectance, $\bar{t}$ and $\bar{r}$ are the means, λ is the number of wavelength channels and α is the similarity coefficient.

To classify the best matching paint, we then take the maximum value between the unknown signature and each signature in the library, as given by Equation 4.13. Naturally, we take the maximum because correlation is highest when the value is closer to 1.

$$Label = \text{argmax}_{\text{i}}(\alpha(t, r_i)) \tag{4.13}$$

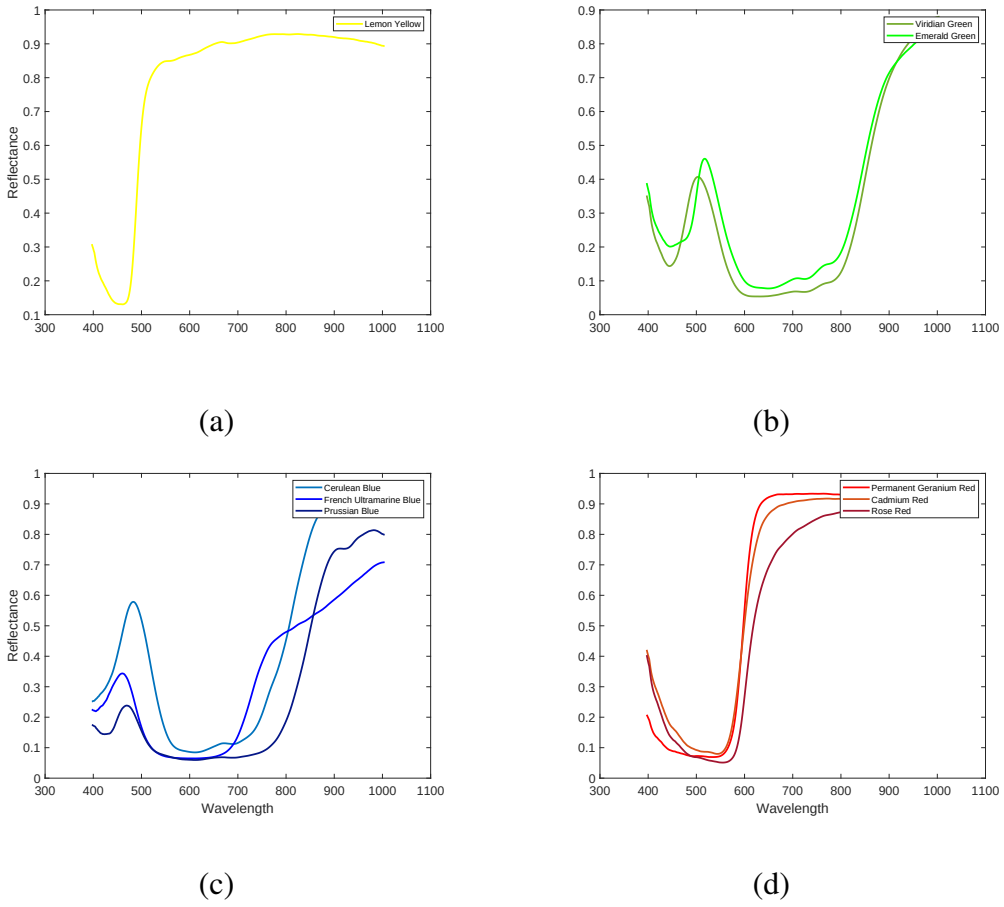

FIGURE 4.17: Pigments spectra separated by their class to show the difference in spectral variance for each class. The Cadmium Red Hue and Permanent Geranium Lake red paints have a very similar signature. On the other hand, although the Viridian Green and Emerald Green paints seem to be very similar, the difference in the location of their reflective peaks is sufficient to allow distinction between the two pigments.

where t and r are the test signature and reference signature, i is the number of paints in the library, $Label$ is the resulting label, and $\alpha(t, r)$ is given by Equation 4.12.

4.5.4.2 Hierarchical NNMF Paint Identification

Once the pure paint regions have been determined, the mixed paint regions must be classified according to the paints in the mixture. As mentioned in the introduction of this section, we use the NNMF algorithm to unmix the features in the mixed signal. In Section 4.5.1, we defined a feature matrix $\boldsymbol{H}$ and a weight matrix $\boldsymbol{W}$ to represent an input data matrix $\boldsymbol{I}$. When our signals are in the form of reflectance, the feature matrix $\boldsymbol{H}$ is of size $\lambda \times k$ where λ are the number variables (wavelengths) and k is the rank, whereas the weight matrix $\boldsymbol{W}$ is of size $k \times N$ where N is the number of data points.

We have also mentioned that we will be using the absorption a and scattering s coefficients, to compare with r. Therefore, after learning the a and s coefficients by nonlinear least squares, we can replace the reflectance r signatures by concatenating a and s. Therefore, $\boldsymbol{H}$ is now $2\lambda \times k$, while $\boldsymbol{W}$ remains unchanged. However, for this to work, the data input matrix must also be in the form of a and s, rather than r.

The metric that has been developed to quantify the similarity between a combination of two pigments and the input data includes both correlation and reconstruction error. The correlation gives us a quantification of similarity between the spectral signatures, whereas the reconstruction error quantifies the difference between the reconstructed result and the input. We normalise the reconstruction error between all combinations such that it will not bias correlation at large values, as given by Equation 4.15. This form of distance metric is similar to Spectral Similarity Scale (SSS) [56].

$$R = H_{2\lambda-1,r} \times W_{r,N} \tag{4.14}$$

$$D = \sqrt{E(I,R)^2 + (1 - SCM(I,R))^2} \tag{4.15}$$

where E denotes the reconstruction error given by Equation 4.4, SCM denotes the Spectral Correlation Mapper given by Equation 3.6, and I and R are the input data and the reconstructed data, respectively. When calculating the reconstruction error, the sum-to-one constraint term is not required, and therefore it is not be used, as Equation 4.14.

To find the paints from each class that best match the mixture, the same approach is applied, in which each combination is tested for reconstruction error. A final result shows the combination of pigments ranked by reconstruction error and shows the weighting of each pigment to create the mixture in $\boldsymbol{I}$.

4.5.5 Method 2 - Global NNMF Paint Analysis

The Global NNMF Paint Analysis utilises full features to estimate weights for each region in the image, with respect to the features presented. The features in this case can either be the paint classes that represent all of our nine paints, or the nine paints themselves individually. In the previous subsection, we used NNMF to find the best matching paints in mixtures without using prior knowledge about the mixture, except assuming that the mixture is made up of two paints. Moreover, the analysis was performed on homogeneous regions, and hence the input data was homogeneous as well. This means that we performed a local optimisation problem rather than a global optimisation problem for all regions in the image. In this subsection, we aim to apply NNMF to a

global optimisation problem for every pixel or region in the image, without having the prior knowledge of any existing mixed regions. By utilising both class features and individual features, we can compare the performance of NNMF with regards to accuracy for correctly classifying both individual and mixed paint regions.

In [20] it is stated that acrylic and oil paints can be considered to be opaque, however we notice that the opaqueness of oil paints is very dependent on the thickness of the layer of paint. In fact, in our mock-up samples, it is noticed that at some areas, the substrate can be seen through the layer of paint. This means that the layer is thin enough such that light is reaching the substrate, and scattered back. In fact, we observe that the concentration of certain paints (eg. Viridian Green Hue) is very diluted, and opaqueness is hard to achieve. Therefore, the reflective signature at these areas will have the same shape, but with a higher amplitude because the white coating of the canvas is equally reflective at all wavelengths in the visible region. Similarly, when these pigments are layered on top of each other, if the top layer is not thick enough, scattered light from the pigment on the bottom may leak through the top layer. By global optimisation, NNMF should be able to identify these areas when the weights $\boldsymbol{W}$ are being learned, therefore, being able to illustrate the membership of each region or pixel with respect to the reference features.

It is important to note that this work has already been attempted by [20]. However, in [20], the authors only addressed the problem of finding different classes of paints in an image. In this work, we also test for the ability to identify specific paints. Therefore, since we have both the pure paints' spectral reflectance, and the average paints' spectral reflectance (of each class), we can compare the accuracy of performing matrix factorisation by NNMF on both features. This means that if $\boldsymbol{H}$ is the basis matrix and $\boldsymbol{W}$ is the weight matrix, by keeping $\boldsymbol{H}$ fixed we can obtain weight coefficients either based on class specificity or on paint specificity. Moreover, in the previous subsection, we were able to learn the absorption and scattering coefficients for our known reflections. One must note that this was being done for a randomly selected number of pixels dependent on the size of the region. In this case, since we do not know that regions exist, or where they are, we would need to measure the absorption and scattering coefficients for each pixel. Consequently, an image with the same resolution but double the wavelengths will be obtained.

The mathematical model applied in this setup is very similar to that applied in the hierarchical approach in the previous subsection. An input data matrix $\boldsymbol{I}$ is to be factorised into a weight matrix $\boldsymbol{W}$ and a feature matrix $\boldsymbol{H}$, and optimised by Equation 4.4. Just as in the hierarchical approach, $\boldsymbol{H}$ is known. Hence, we only optimise for $\boldsymbol{W}$ while

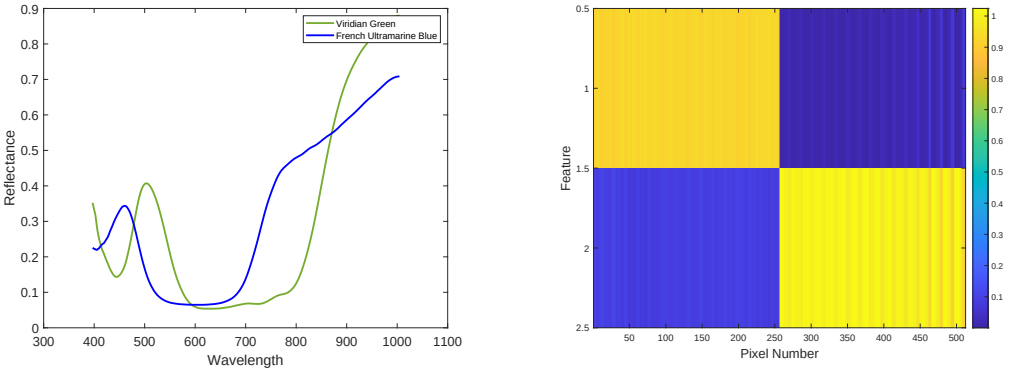

(a) (b)

FIGURE 4.18: Illustrating the features and weights obtained after performing NNMF for a rank of 2. (a) Plot showing the feature vectors used in the NNMF basis matrix. These correspond to French Ultramarine Blue and Viridian Green paints (b) the weight matrix obtained by the NNMF for pixels of which the first 256 correspond to paint feature 1 and the remaining pixels correspond to paint feature 2. As expected the NNMF weight matrix returns a weight value of 1 for the corresponding feature and a weight value of 0 for the other, non-corresponding feature.

keeping $\boldsymbol{H}$ fixed. This means that, the vertices (features) in NNMF space are known prior to calculating the weights. Moreover, by the sum-to-one constraint [57] a signature that fits very closely to either of the vertices should have a weight close to 1, whereas if the signature is the complete opposite it should have a value close to 0, as shown in Figure 4.18. Additionally, the diagram in Figure 4.19 illustrates how the weights could be distributed for a non-negative matrix factorisation of rank 3, in which the features include Viridian Green, French Ultramarine Blue, and the substrate. In the previous subsection, the rank k was always set to two, because we assumed that the maximum number of different paints in a mixture is two, and those regions that contained only one paint were identified prior to the unmixing stage. For this method, k is set according to the maximum number of different paint classes or paints. That is, if we are identifying classes, k will be set to the number of different classes. Similarly, for different paints, k will be set according to the number of different paints in the image.

4.5.6 Method 3 - Direct Classification of Mixtures

When mixing paints, it is very hard to replicate the same mixture for different samples, especially when not using a weighing scales. It is evident, that artists never mix their paints using a weighing scales, because it is time consuming and they rely on intuition. In this method, we will investigate the accuracy of the Kernel Spectral Correlation Mapper

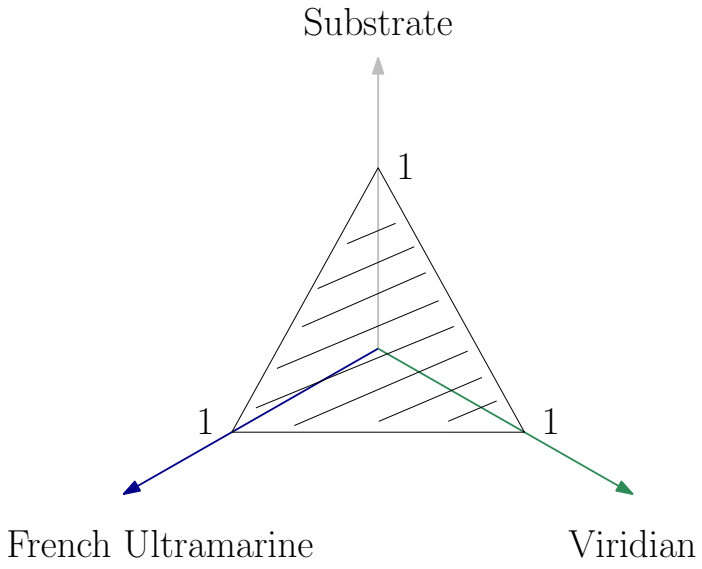

FIGURE 4.19: An example to show how the weights can be distributed by an NNMF of rank 3. The three features are the substrate (in grey), French Ultramarine (in blue) and Viridian (in green). The vertices of the triangle correspond to instances where a single feature has a weight of 1 while the remaining two features have 0 weight indicating the pure form of a feature. Any other point in the area enclosed by the vertices represent a mixture of the three features, which sum to one.

(KSCM) metric and NNMF algorithm when it comes to identifying mixtures by looking for a best mixture match in a reference dataset.

In this subsection, mixtures are treated as individual paints. This means that in this method, we will not be looking to identify the paints making up the mixtures, but rather aim to classify the mixtures. Therefore, we will need to extend our reference library and include a reference signature for each mixture mentioned in Section 2.2. Since the mixtures are now included in the reference library, the method can be a simple correlation method using the previously mentioned KSCM or NNMF, and no hierarchical approaches are required.

For the case of the KSCM the same approach as in Section 4.5.4.1 is applied, following Equation 4.13, where the maximum correlation is accepted as the best matching reference paint. Similarly for NNMF, the same approach as in Section 4.5.5 is applied. The difference is that in this method, the rank k will be larger than that of the Global NNMF Paint Analysis method, because the library also includes the mixtures.

4.5.7 Sensitivity to Variance of Spectra

From literature [18] we know that algorithms, such as NNMF, are better than a FastICA algorithm when the correct ratio of paints in a mixture must be identified. What we do not know, however, is whether NNMF and KSCM are robust enough to accurately identify mixed paint variants when the individual components are not known. In other words, the robustness of both algorithms will be analysed when the input data intra-class variance is large.

Additionally, from samples that we prepared over a range of paints, a difference of 20% in mass to mass ratio can result in a clear difference in the mixed paint. However, the components in the mixture are still the same.

The created variants of each mixture range from mixtures of 60:40 and 40:60 mass to mass ratios, which will be checked against the reference 50:50 mass to mass ratios. The reason why we created these variants is because we want to test the robustness of both classifiers, and to see which classifier is better in terms of accuracy.

4.6 Summary

In this methodology section, a series of methods regarding segmentation of paintings and paint analysis, for the purpose of improving the analysis of paintings, were proposed. We also note down some steps taken to create suitable samples and data such that proper testing can be carried out, and important pre-processing steps required to transform and clean the data.

The segmentation method focuses on keeping a high separability between different paints, while keeping the boundary adherence high as well. Most importantly, our aim was to reduce the over-segmentation caused by existing methods, as these would hinder or slow down the analysis process. To be able to do this, we proposed a new metric that is based on SCM and KSAM, in which the kernel properties of KSAM are applied to the correlation from SCM.

Regarding classification, we proposed three different methods that could be useful in the analysis of these regions extracted from a painting. Existing methods, such as [18] show that an accurate ratio of each paint in a mixture of two paints can be estimated. This, however, is based on prior knowledge regarding the mixture. That is, both paints are known beforehand. For the Hierarchical Paint Analysis and Global NNMF Paint Analysis, we do not make use of such prior knowledge. In fact, these two algorithms must find the paints that create the mixture of paints. For comparative purposes, we

created a third method, Direct Classification of Mixtures, which applies the metrics to the problem of classifying mixtures as individual paints.

Chapter 5

Results and Discussion

In this chapter, the results achieved by our proposed segmentation method will be presented and compared to other existing methods in the paint analysis field. We first evaluate the results obtained by the proposed segmentation method. Then we evaluate and compare the results obtained by the classification methods. Finally, we show the results for the proposed segmentation algorithm using 'The Birth of John the Baptist' by Niccolò da Bologna, as a case study.

5.1 Segmentation Evaluation Methodology

The evaluation for the proposed segmentation algorithm comprises of both a qualitative and a quantitative study. The qualitative study entails a comparison against other state-of-the-art algorithms, and additional tests performed to test the capabilities of the algorithm. On the other hand, the quantitative study entails a comparative study using specific metrics against the same state-of-the-art algorithms. The metrics used for the quantitative study are described in Section 5.1.2.

5.1.1 Qualitative Evaluation Methodology

The proposed, Spectral Similarity Merging (SSM), consists of a three-stage region-merging process, which aims to maximise homogeneity and boundary adherence. Image segmentation algorithms are assessed on their ability to segment different regions, while keeping homogeneous regions whole. One other important aspect is that the detected region boundaries are located as closely as possible to the true region boundaries. This section describes and discusses the qualitative analysis of the proposed SSM, and compares the performance against other state-of-the-art segmentation algorithms.

5.1.1.1 Region Merging

The first stage of the algorithm consists of the merging of initial square regions that have been seeded over the image. Each region, represented by a spectral signature, is compared with neighbouring regions for merging. Ideal results for this stage would show that regions that have a similar spectral signature are merged into a single region. Since we do not perform any additional processing at this stage, we expect regions that are not homogeneous to remain un-merged. Since the Kernel Spectral Correlation Mapper (KSCM) algorithm is based on correlation, we expect regions that have a similar spectral signature, but have a difference in amplitude, to still merge as a single region. We test the output of the region merging stage by visual inspection for the following criteria:

- Homogeneous regions are merged with similar neighbouring regions.
- Non-homogeneous regions remain un-merged.
- No paint regions leak into other paint regions.

5.1.1.2 Pixel Re-Assignment

In this subsection, a similar qualitative analysis is carried out. This section includes two sub-algorithms as mentioned in Section 4.4.1.2. The first algorithm deals with high intra-class variance regions, whereas the second deals with low intra-class variance regions. Regions with high intra-class variance are analysed for the possibility of having multiple paint signatures. Ideal results for this stage would show regions containing multiple paint features to be split and the different features merged with similar neighbours. Moreover, we should expect no fragmentation of these regions because the re-assignment of pixels is controlled by a spatial term. To demonstrate the importance of this spatial term, a test is carried out in which no spatial term is included such that the results can be compared. For this sub-stage, the results will be evaluated against the following criteria:

- Re-assignment of high intra-class variance regions do not produce ‘orphaned’ pixels.
- Re-assignment of high intra-class variance regions produce a smooth boundary of where the two paints meet.
- No paint regions leak into other paint regions.

Regions with low intra-class variance are analysed for the smoothness of the boundary neighbouring other low intra-class variance regions. Ideal results for this stage would

show that the boundary of a low intra-class variance regions adheres to soft edges (those edges which are created by a difference in layer thickness). We visually inspect the output for these stages based on the below criteria:

- Re-assignment of low intra-class variance regions do not produce 'orphaned' pixels.
- Re-assignment of low intra-class variance regions produce a smooth boundary of the 'weak' edges in the homogeneous region.
- No paint regions leak into other paint regions.

5.1.1.3 Checking for Further Merging

Since this algorithm is very similar to the Region Merging stage, the test procedure is also similar. However, since we have processed our regions in the Pixel Re-assignment stage, we expect to see additional merging. The output for this stage is visually inspected for the following criteria:

- The number of regions when compared to the Pixel Re-assignment stage has reduced.
- No paint regions leak into other paint regions.

5.1.1.4 Complete Algorithm

To demonstrate the improvement that SSM brings to the segmentation of paintings, we evaluate the result of the complete algorithm and directly compare with other state-of-the-art algorithms. We also compare the hyperspectral domain against the RGB domain, to show that the hyperspectral domain presents additional features that are crucial in separating very similar regions. The specific algorithms that will be used for evaluation are a Sobel Edge detector, a K-means algorithm [61], a Simple Linear Iterative Clustering (SLIC) superpixels algorithm, and SSM. The results will be evaluated against the following criteria:

- Minimal, or no under-segmentation of regions.
- An improvement in the amount of over-segmentation of regions.
- An improvement in F1-score (includes recall and precision).
- Robustness improvement over the under-segmentation/over-segmentation when varying the starting number of superpixels.

5.1.2 Quantitative Evaluation

In the previous sub-section, we have described several qualitative tests to investigate whether SSM performs better than other state-of-the-art algorithms. To demonstrate this even further, in this sub-section, we apply well-known segmentation metrics, such as over-segmentation, under-segmentation, and F1-score, to quantitatively evaluate the improvements of SSM over the other algorithms. The ground truth image for these metrics is an edge pixel image of the boundaries of regions. Similarly, the segmentation image is transformed into an edge pixel image such that they can be compared.

5.1.2.1 Measuring Over-Segmentation and Under-Segmentation

One important metric in this study is over-segmentation. In our study, we define over-segmentation to be the ratio of the amount of detected edge pixels to the ground truth edge pixels. The detected edge pixels that correctly identify a ground truth edge pixel, are accepted as true positives TP_L, while those that do not are considered false positives [45]. The over-segmentation value is measured by Equation 5.1:

$$OS = 1 - \frac{|TP_L|}{|L|} \tag{5.1}$$

where $|TP_L|$ is the cardinality of true positive pixels from the detected edges, $|L|$ is the cardinality for all edge pixels in the detected edges, and OS is the over-segmentation result in which the ratio for false positive pixels in the detected edges is given. The tolerance set for accepting true positive pixels is of 9 pixels, such that ground truth edge pixels that may have been badly perceived in the ground truth do not affect the score of the metric. This tolerance is also applied to the under-segmentation and F1-score metrics.

We define under-segmentation to be those pixels that are in the ground truth edge pixel image, but are missing in the detected edge pixel image, TP_G. Hence, under-segmentation is the measure of the amount of false negatives in the detected edge pixel image [45]. This measure is defined by Equation 5.2:

$$US = 1 - \frac{|TP_G|}{|G|} \tag{5.2}$$

where $|TP_G|$ denotes the cardinality of false negatives from the detected result, $|G|$ is the cardinality for all edge pixels, and US is the under-segmentation result, which gives the ratio of all ground truth pixels in comparison to missing pixels in the detected edge result.

5.1.2.2 Measuring F1-Score

Similarly, the F1-score is measured using these edge pixel images. F1-score is the harmonic mean between precision and recall. Precision is the ratio of the amount of true positives against false positives whereas recall is the ratio of the amount of true positives against false negatives. The F1-score is defined by Equation 5.3:

$$F = \frac{2 \times R \times P}{R + P} \tag{5.3}$$

where R is the Recall, P is the Precision, and F denotes the F1-score [62].

5.1.2.3 Measuring Detected Pixel Offset

The previously mentioned metrics describe the amount of detected edge pixels that are accepted as true positives, false positives, or false negatives. However, these measures do not describe the amount of offset that our segmentation has with respect to the ground truth. Therefore, we are also including a pixel offset metric, which measures the pixel distance between the detected edge pixel image and the ground truth edge pixel image. Each pixel in the detected edge pixel image finds its closest ground truth pixel and saves that distance. This metric is defined in Equation 5.4:

$$\boldsymbol{d} = \mathrm{argmin_i}(E(L_i, G)) \tag{5.4}$$

where E denotes Euclidean Distance, $\boldsymbol{d}$ contains all the distances for the detected pixels, L are the detected pixels, and G are the ground truth pixels.

5.1.2.4 Sensitivity to Initial Parameters

One final evaluation procedure that will be carried out aims to test the sensitivity of the algorithm to changes in initial parameters. The initial parameter for the starting number of superpixels, k, will be varied across each sample image and algorithm. This will be done to demonstrate the sensitivity of each algorithm toward this parameter, and observe whether the resulting metrics worsen or improve in terms of the metrics discussed earlier.

The initial parameters will be varied as follows: For images 1,2 and 4, the number of superpixels will be varied from 5,000 to 15,000 with 10,000 being the control; For Image 3 the number of superpixels will be varied from 500 to 3,500 superpixels with 2,000 being the control. The reason why Image 3 has been assigned less initial superpixels is because we have manually measured the appropriate size of superpixels such that the smallest region can be sampled.

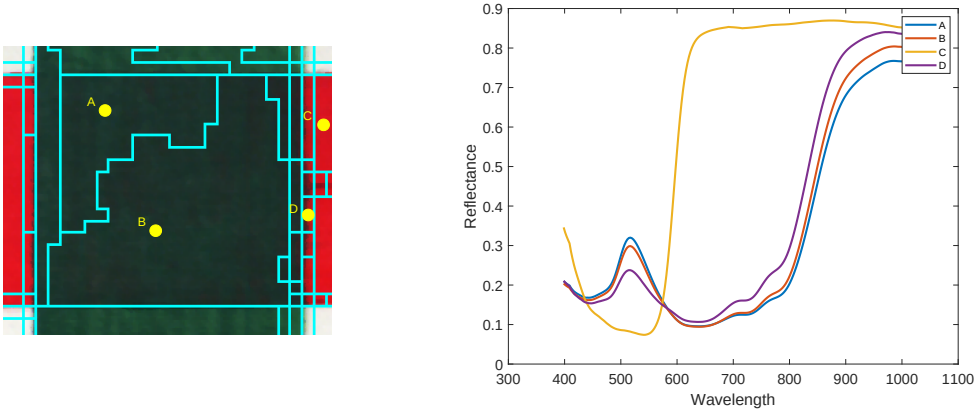

(a) (b)

FIGURE 5.1: The result from the Region Merging stage shows that the initial criteria are obeyed but with room for improvement. In (a) initial regions have been merged according to the spectral signature shape (b) shows the plots for the spectral centroid of the regions marked in yellow.

5.1.3 Qualitative Evaluation Results and Discussion

In this section, the results from each stage of the algorithm are illustrated and evaluated qualitatively. A discussion on each stage is also given, focusing on the improvements and points of failure of each stage, as well as demonstrating the importance of certain measures taken when designing the algorithm. For the stage-by-stage evaluation, we focus on a cropped region from Image 1 to highlight the outcomes. The cropped region features two paints which overlap each other at the central square. The paints are Cadmium Red Hue and Viridian Green Hue, where the Viridian Green Hue is the paint overlapping the Cadmium Red Hue.

5.1.3.1 Region Merging Results

Figure 5.1 demonstrates the performance of the Region Merging stage of the algorithm, showing that the proposed method separates different paints while grouping homogeneous superpixels together. Once grouped together, some of these homogeneous regions stop growing, keeping some over-segmentation. When closely analysing the spectra of these un-merged regions, we notice that although KSCM has some degree of tolerance towards change in amplitude, the sensitivity of the radial basis function in use decreases the degree of correlation.

In other regions, especially those where the Viridian Green transitions from an individual paint area to an area where the Viridian Green overlaps the Cadmium Red,

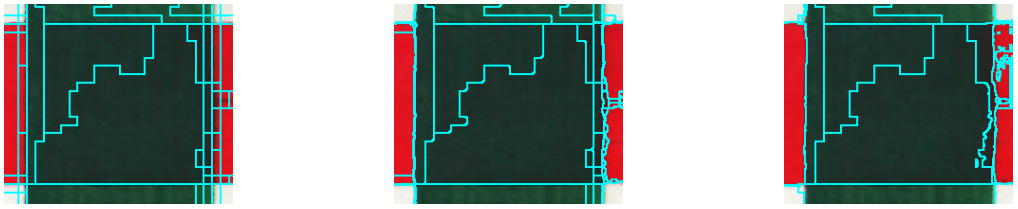

(a) (b) (c)

FIGURE 5.2: The result from Pixel Re-Assignment sub-routine 1 shows that the boundary adherence to true edges is strong after the re-assignment of high intra-class spectral variance regions (a) shows the input image to this stage where the regions are separated by the cyan lines, (b) shows the result for Pixel Re-assignment sub-routine 1, and (c) shows the result if the spatial term is ignored for the same method.

the algorithm fails to detect a high variance in the spectra. Hence, this results in missed high intra-class variance superpixels. Additionally, this may also be caused due to the chance that a superpixel boundary corresponds exactly over the true edge of two adjacent regions.

The overall goal for this stage was to improve on the reduction of over-segmentation without introducing under-segmented regions. The result discussed in this section indicates that this goal has been well achieved.

5.1.3.2 Pixel Re-Assignment Results

The second stage of the algorithm contains two sub-routines. The first sub-routine deals with the processing of high intra-class variance regions. The second sub-routine deals with the processing of low intra-class variance regions. The results for the first sub-routine are illustrated in Figure 5.2, whereas the results for the second sub-routine are illustrated in Figure 5.3.

From the results shown in Figure 5.2, we demonstrate the pixel re-assignment of the high variance regions, as well as the importance of the spatial term. The spatial term allows some control of the direction in which a pixel from a high intra-class variance region moves to. Without the spatial term, one can observe from Figure 5.2(c) that when correlation is only marginally different for different superpixels, pixellation from the re-assignment occurs. Since this sub-routine is iterative, it also follows a variance minimisation criteria, until all regions are below the chosen variance threshold. The goal for this stage was to re-assign pixels in high intra-class variance regions, based on their

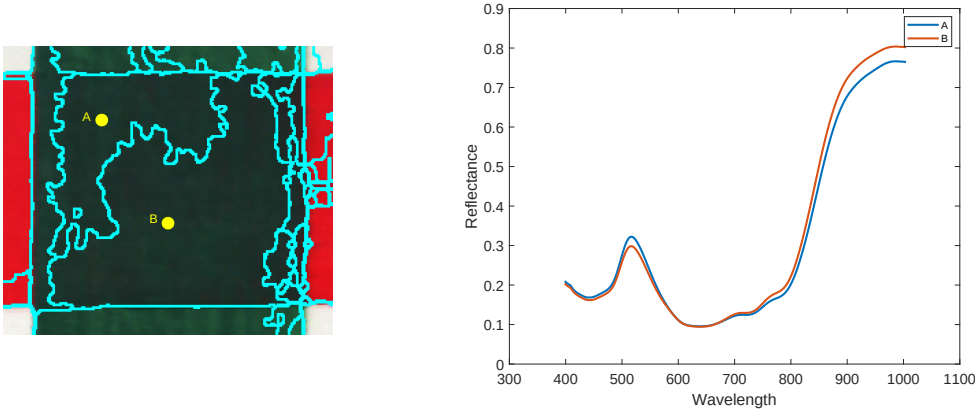

(a) (b)

FIGURE 5.3: The result from Pixel Re-Assignment sub-routine 2 shows that the boundary adherence to soft edges is strong after the perimeter smoothing of low intra-class spectral variance regions (a) shows the result for Pixel Re-Assignment sub-routine 2 (b) shows the similarity between adjacent regions of the same hue.

correlation with neighbouring regions but controlled with a spatial term. This goal is achieved by this sub-routine, increasing the homogeneity of superpixels.

Figure 5.3 demonstrates the results obtained by the second sub-routine. The result for this stage is satisfactory, because the superpixel borders are refined based on correlation to any neighbouring region, without introducing new under-segmentation. Moreover, by having a smoother correlation-based perimeter, the homogeneity of each superpixel is further increased.

5.1.3.3 Checking for Further Merging Results

This stage is fundamental because it allows homogeneous regions to merge after pixel re-assignment sub-routines. As mentioned, the Pixel Re-Assignment stage provides stronger homogeniety in our superpixels. Therefore, we re-use the Region Merging stage, to check for new highly correlated neighbouring regions. In fact, in Figure 5.4 we show that by re-applying the Region Merging routine, we can mitigate some of the over-segmentation.

In some cases, we observe that nested regions were developed. In our samples, these regions were usually made of the same paint as their surrounding region. However, this may not always be the case for real scenarios, since, it is possible to paint in a concentric circle manner. Therefore we cannot just merge these regions with their surrounding neighbour. However, we can still make use of a relaxed threshold because for different paints the correlation from KSCM should be significantly small.

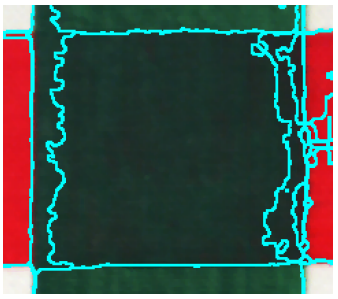

FIGURE 5.4: The result from stage 3 shows that the over-segmentation is further reduced without introducing additional under-segmentation.

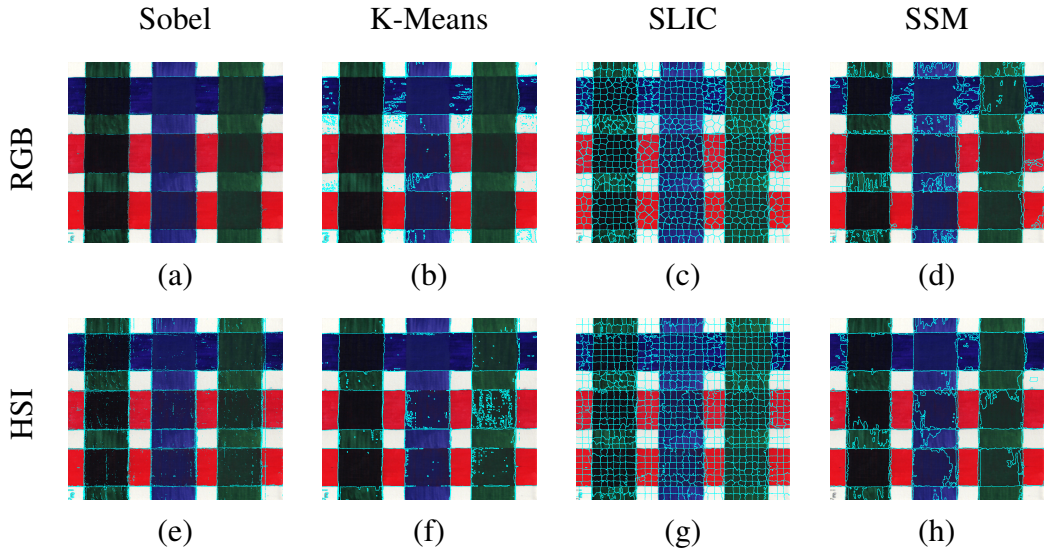

FIGURE 5.5: Comparing different segmentation techniques for the task of paint segmentation. Rows represent the results for the type of data: Row 1 are results obtained from an RGB representation of the hyperspectral image, and Row 2 are results obtained for the hyperspectral image. Columns represent the algorithm: column 1 is a Sobel Edge detector, column 2 is K-means clustering, column 3 is SLIC superpixels, and Column 4 is the proposed SSM.

Following this stage, a post-processing step was applied to remove small regions made up of a single pixel, because we assumed that it is highly improbable for an artist to paint - using brushes - single pixel large features.

5.1.3.4 Complete Algorithm Results

In the previous subsections, we have seen the results acquired by each stage of SSM. We have seen that the criteria have been met, resulting in successful segmentation of our mock-up samples. In fact, in Figure 5.5 we demonstrate the results obtained by multiple segmentation methods, illustrating the effectiveness of our proposed method. One can

observe from the figure that over-segmentation is much less than for the SLIC, for the same number of initial superpixels. Although SLIC does not under-segment the image, it does this at the cost of over-segmentation. Additionally, much less under-segmentation is achieved by the proposed method when compared to the Sobel Edge detector and the K-means clustering algorithm.

Figure 5.5 also demonstrates the difference between processing an RGB domain image and a hyperspectral domain image, illustrating the effectiveness of using hyperspectral features over RGB features. A clear improvement is observed when paint regions having similar hues are adjacent to each other, because the difference in features is clearer in the hyperspectral domain. This, however, produces a segmentation improvement for all processing methods, including the Sobel Edge detector, K-means clustering, SLIC and SSM. Therefore, we can conclude that SSM is affected in a similar way to the other methods when the number of features per pixel vary.

Figure 5.6 illustrates the results obtained for each mock-up sample image, highlighting the obtained region edges against the ground truth region edges. Since the Sobel Edge detector does not separate the regions according to spectral homogeniety, but rather according to change in intensity, we do not include the Sobel Edge detector for the results shown in Figure 5.6. Similarly, the Sobel edge detector is not included in the quantitative evaluation discussed in the next section.

The results obtained demonstrate that the proposed SSM accurately separates homogeneous regions by their spectral signatures. Moreover, it also demonstrates the capability to separate regions that have a very similar hue, specifically adjacent red paints of similar shades. However, a special case in Image 4 can be observed, in which two adjacent regions are merged as one region. Figure 5.7 focuses on this case, illustrating that this merging occurs due to the fact that the paints in the two mixtures are the same, although the paint mixtures in two regions are of different ratios. Therefore, the change in correlation is too small to be captured by SSM. In fact, the graph in Figure 5.7 shows that the spectral curve of both regions is basically the same. Similarly for SLIC, the results suggest that both regions cannot be properly distinguished.

Recalling the fact that the proposed algorithm is based on using the optimal threshold, and the chosen threshold is based on all four images, we tested the possibility of using a stricter threshold to check whether the two paint mixtures can be separated. This test is demonstrated in Figure 5.8, illustrating the results from different thresholds. The result indicates that by using a stricter threshold, it is possible to separate the two highly correlated regions, at the cost of slightly higher over-segmentation. The image under test is a cropped version of Image 4, highlighting the region of interest. The initial

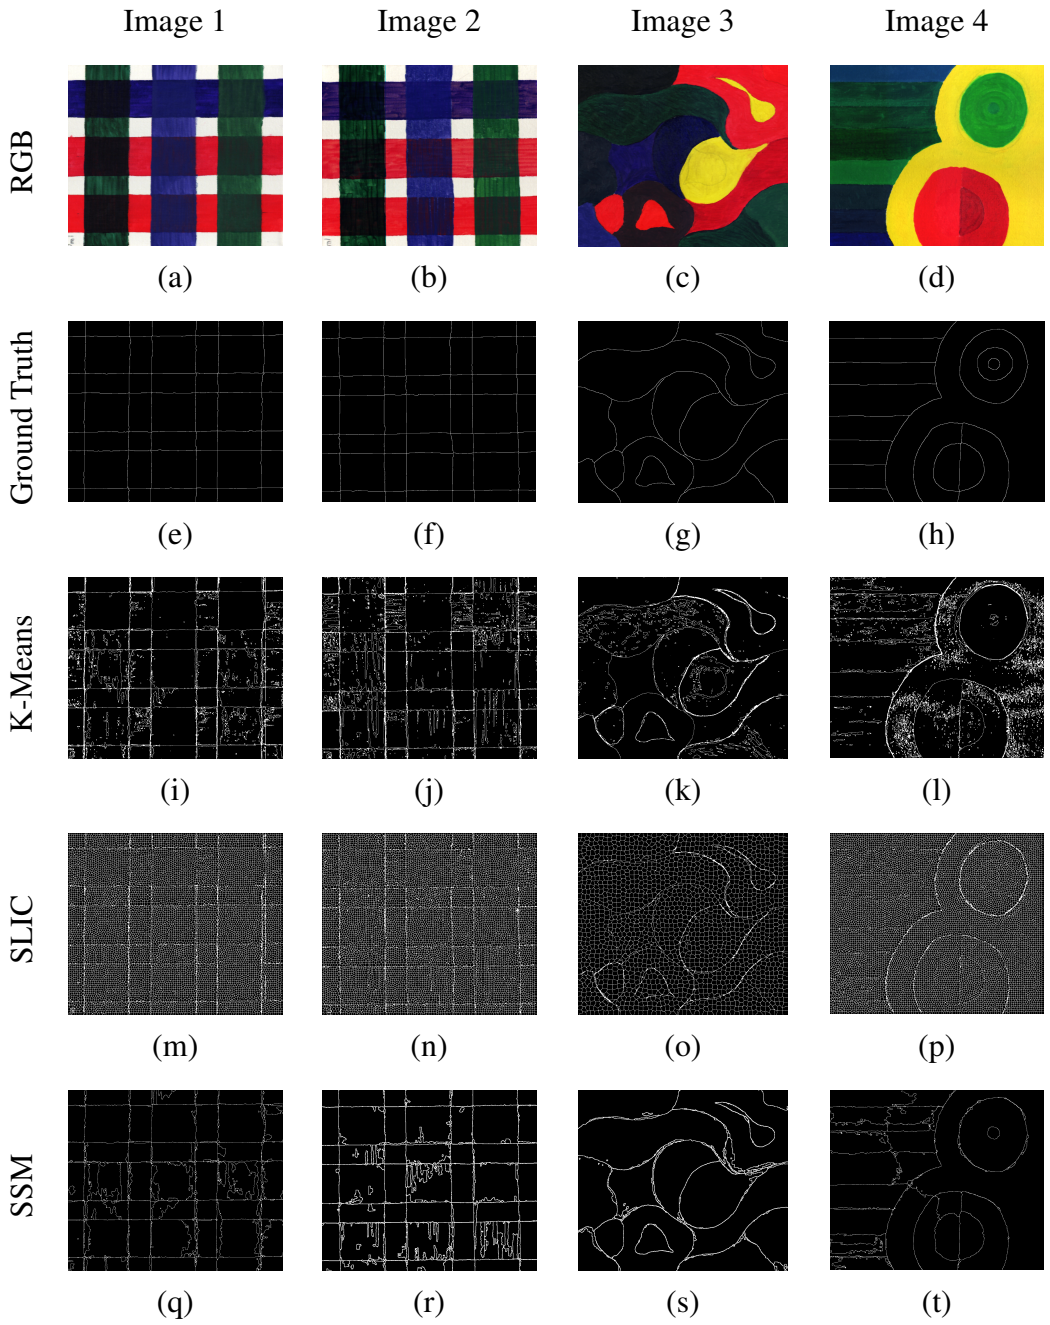

FIGURE 5.6: Results for four different samples images, separated by columns. Row 1 represent an RGB representation of the HSI image, Row 2 represent the ground truth image, Row 3 represent the result from K-means (k = 16,16,12,15), Row 4 represent results for SLIC (k = 10000,10000,2000,10000) and Row 5 represent results for SSM (k = 10000,10000,2000,10000).

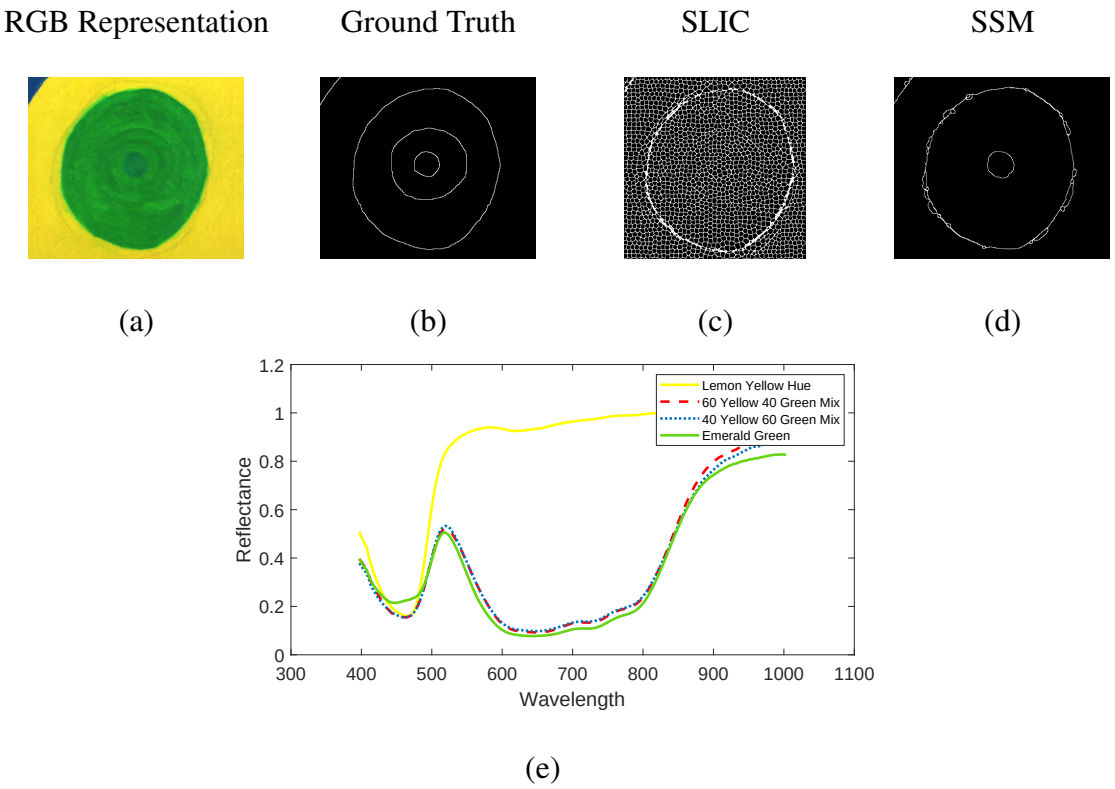

FIGURE 5.7: A case where paint separation fails: (a) the RGB representation of region that fails to merge, (b) the ground truth for (a), (c) the resulting segmentation for SLIC, (d) the resulting segmentation for SSM, and (e) a graph showing the spectral curves for the 4 regions shown in the image.

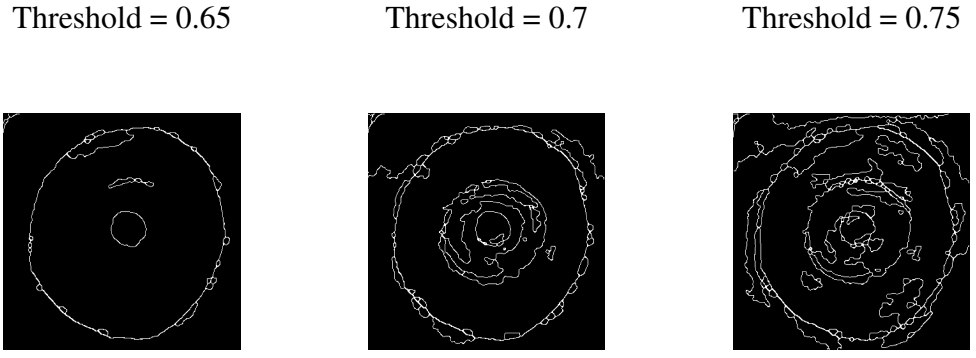

FIGURE 5.8: Varying the initial region merging threshold T_1 to (a) 0.65, (b) 0.7, and (c) 0.75 to test whether we are able to capture the inner region if the threshold is increased.

parameters are kept constant, that is $k = 10,000$ and the threshold increased from 0.65 to 0.7 and 0.75. For these parameters, an F1-score of 0.47, 0.37 and 0.25 is acquired, for the respective thresholds.

5.1.4 Quantitative Evaluation Results and Discussion

The same image data sets will be used for each of the metrics mentioned in the following sub-sections. To be able to compare the K-means with SLIC and SSM, an optimal number of k clusters was found. The optimal number of k clusters for each sample image is as follows: $k = 16$ for Images 1 and 2, $k = 12$ for Image 3 and $k = 15$ for Image 4. On the other hand, the number of initial superpixels used for both SLIC and SSM is as follows: $k = 10,000$ for Images 1,2 and 3, and $k = 2,000$ for Image 3. As mentioned in Chapter 4, we measure the initial number of superpixels according to the size of the smallest region in the image. Although this value satisfies the Nyquist criteria, it does not mean that we cannot over sample. In fact, over-sampling can help - although not within the scope of this study - in SSM because it can increase the ability to capture finer details.

In the following sub-sections, we show the results obtained for each metric and discuss the main differences for the specific metric between the three algorithms used.

5.1.4.1 Over-Segmentation and Under-Segmentation Results

The results in Table 5.1 show that SSM reduces the over-segmentation by 27.5% when compared to SLIC. Although the over-segmentation for the K-means algorithm is smaller when compared SLIC, it is still much higher than SSM.

Although we carefully chose k clusters to match the number of clusters in the image, there were still some 'noise' clusters assigned. This resulted in increased pixellation and over-segmentation, therefore reduced precision. When compared to SSM, the over-segmentation of the K-means is 19.75% higher.

Table 5.1 also shows that an increase in under-segmentation is obtained when SSM is compared to SLIC. Although this is not desired, it may be due to the fact that there may be some ambiguity between the true edges of regions and the ground truth edges of regions. In fact, the increase of under-segmentation is that of 0.075% with both methods having an average of lower than 1% total under-segmentation.

On the other hand, when comparing SSM to K-means, we observe that the K-means has an additional 0.125% under-segmentation over SSM. Therefore, it could be said that the under-segmentation produced by each method is significantly small.

5.1.4.2 F1-Score Results

The F1-Score measures the Harmonic mean between Precision and Recall, recorded in Table 5.2. The results demonstrate that the F1-score improves when comparing SSM

TABLE 5.1: Over-Segmentation (OS) and Under-Segmentation (US) results for SSM when compared to K-means and SLIC.

	Over-Segmentation (%)			Under-Segmentation (%)		
	KMeans	SLIC	SSM	KMeans	SLIC	SSM
Image 1	85	95	68	0.2	0	0.1
Image 2	85	95	70	0.2	0	0.1
Image 3	86	93	67	0.3	0.1	0.1
Image 4	92	96	64	0.2	0	0.1

TABLE 5.2: F1-Score results for SSM when compared to K-means and SLIC.

	Precision			Recall			F1-Score		
	KMeans	SLIC	SSM	KMeans	SLIC	SSM	KMeans	SLIC	SSM
Image 1	15	4	32	100	100	99	26	8	48
Image 2	14	4	30	100	100	99	24	8	46
Image 3	14	7	33	96	99	98	23	13	50
Image 4	8	3	36	74	100	85	15	6	50

with both SLIC and K-means, showing a greater improvement over SLIC. However, if we take a closer look at the precision and recall, we observe that SLIC suffers greatly from low precision due to the high over-segmentation that it produces. On the other hand, because of the low under-segmentation, the recall measure is very high since a very low amount of false negatives would be obtained.

5.1.4.3 Detected Pixel Offset Results

We used the under-segmentation and over-segmentation metrics to measure the ratio of true positives to both the false negatives and false positives, respectively. These results are measured according to the tolerance of 9 pixels that we allowed when comparing the detected pixel edge images with the ground truth edge images. However, these results do not tell us how much the detected edge pixels are offset from the ground truth edge pixels. The histogram presented in Figure 5.9 illustrates the frequency of the distance for the detected pixels from the ground truth pixels. Only one histogram is shown for all four mock-up images. An individual histogram for each mock-up image can be found in Appendix A of this dissertation.

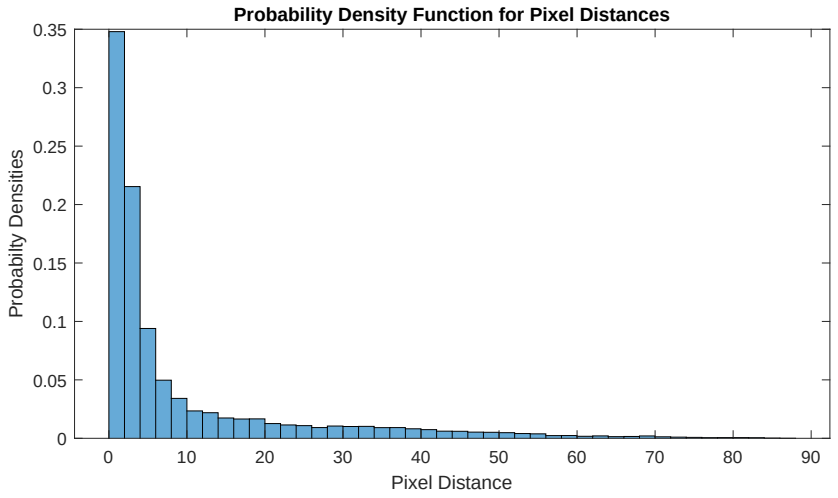

(a)

FIGURE 5.9: Pixel Distance Probability Density plot for all mock-up samples. The plot shows that the majority of detected pixels are very close to the ground truth pixels.

Results show that the majority of the detected pixels are within a distance of 10 pixels, ensuring that the 9 pixel tolerance chosen for the previously mentioned metrics was not allowing ambiguously detected pixel edges to be accepted as true positives.

5.1.4.4 Sensitivity to Initial Parameters Results

A test in which the initial superpixel parameters of 10,000 superpixels for Images 1,2 and 4, and 2,000 for Image 3, are varied between 5,000-15,000, and 500-3,500, respectively, to measure the sensitivity of SSM and SLIC is carried out in this subsection. Results for this test are demonstrated by a box plot, as shown in Figure 5.10. The box plot demonstrates that although SSM is slightly affected by a change in the number of superpixels, the median value is varying much less than for SLIC for both over-segmentation and F1-score.

5.2 Classification Evaluation Methodology

In this section, we will be evaluating the paint classification problem of this dissertation. In particular, we are interested in identifying both individual and mixed paint regions, and try to identify which paints make up the mixtures. The methods described in Chapter 4

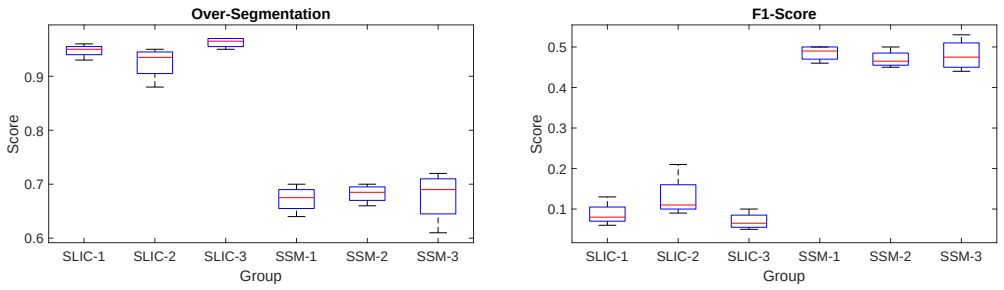

(a) (b)

FIGURE 5.10: A box plot for (a) over-segmentation and (b) F1-score, show the improvements and losses that our algorithms has on the mock up samples we provided.

Section 4.5 will be evaluated for their ability to determine paints in paintings. This section describes the evaluation methodologies and the results obtained for each method.

For the classification evaluation, we focus on quantitative results, to show the differences in accuracy for the three different paint classification methods being investigated. Each method is evaluated separately and a direct comparison between all methods can be found at the end of this section. To our knowledge, identification of individual paints from paint mixtures through HSI has not been reported and thus, we cannot directly compare to other work.

5.2.1 Hierarchical Paint Analysis - Evaluation Methodology

The Hierarchical Paint Analysis method is a two-stage algorithm that uses the results obtained from SSM to assign a paint label to the image pixels. Since the Hierarchical Paint Analysis method is a two-stage algorithm, the evaluation will be separated into two sub-sections. The first stage is the correlation stage, in which KSCM is used to detect regions made up of individual paints, and classify them according to the correlation with a reference library. The second stage is the local optimisation stage, in which Non-negative Matrix Factorisation (NNMF) is used to classify the patches of mixed paints and determine which individual paints make up the mixed paint.

5.2.1.1 Measuring KSCM Accuracy

To measure the accuracy at the correlation stage, we need to measure how many individual paint regions have been correctly classified in terms of individual paint. Additionally, we also need to distinguish mixed paint regions from individual paint regions. Therefore, the aim of KSCM is to identify these individual paint regions from mixed paint regions.

Following these definitions, two separate accuracy measures will be recorded. The first measures the accuracy for the total number of correctly identified individual and mixed paint regions. The second accuracy metric measures the total number of correctly identified paint in each individual paint region.

For this stage only, the KSCM results are compared to similar metrics, such as Kernel Spectral Angle Mapper (KSAM), Spectral Angle Mapper (SAM) and Spectral Correlation Mapper (SCM), to demonstrate the effectiveness of our proposed correlation metric.

5.2.1.2 Measuring Hierarchical NMF Accuracy

For the second stage, the accuracy is measured in a similar manner to the correlation stage, that is, the accuracy metric measures the total number of correctly identified individual paints for each mixed paint region. Since this stage operates in a hierarchical manner, two accuracy measures are computed. The first measure gives the accuracy for correctly identifying the classes of the individual paints in the mixture, whereas the second measure gives the accuracy for correctly identifying the specific paints in the mixture.

5.2.2 Global NNMF Paint Analysis - Evaluation Methodology

The Global NNMF Paint Analysis method is a single-stage algorithm, in which NNMF is used to solve a global optimisation problem to classify the regions obtained from SSM, similar to the Hierarchical method. The goal for this evaluation is to determine the accuracy of the method with regards to the capabilities of classifying both individual paint regions and mixed paint regions at the same time.

5.2.2.1 Measuring Global Optimisation NMF Accuracy

To measure the accuracy for this method, each paint region obtained by SSM is checked for having a high weight coefficient. Since NNMF is a linear algorithm, and the method includes all pixels and features, a high weighting coefficient corresponding to a region, say a red region, may not be as high as a high weighting coefficient in another region, say a green region. However, by viewing each region's weight coefficient for each feature, we determine the highest weighting for that particular feature. Based on the weight coefficient for a single feature, we then determine if the region is classified correctly. A true positive (TP) for this metric is considered when a region containing the highest weighting corresponds with the ground truth. If the highest weight is present in a feature

that does not correspond to the ground truth, then this is recorded as a false positive (FP). Similarly, true negatives (TN) would be those regions which are not detected in a feature that does not represent them, while false negatives (FN) are those regions which are not detected in features that represent them. Equation 5.5 is used to measure the accuracy metric.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{5.5}$$

In this method, the possibility to determine paint regions by having multiple paints of the same class rather than a single paint from one paint class is demonstrated using NNMF. The possibility to determine specific individual paints in mixed regions by studying the weight coefficients of such regions is also demonstrated.

5.2.3 Direct Classification of Mixtures - Evaluation Methodology

The Direct Classification of Paint Mixtures classifies the mixed paint inputs against a mixed paint reference dataset. The aim is to observe whether including paint mixtures in reference datasets increases the accuracy of mixture identification.

5.2.3.1 Measuring Mixed Paint Reference Library Accuracy

The main difference of this method from the Hierarchical method and the Global Optimisation method is that the reference library contains mixed paint references. Therefore, one of the goals of this method is to determine and observe accuracy changes in classification of these paint regions by including mixtures in reference libraries. Having mixture reference spectra will increase the chances that a mixed paint region is classified correctly. However, some mixtures created from different paints have very similar hues and spectral curves. We aim to demonstrate that although these mixed paints have similar spectral curves, both KSCM and NNMF are capable of correctly classifying the mixed paints. For this algorithm, we prepared a mock-up sample containing swatches of paints. We use these mixed paints to create the reference library by extracting data from each swatch and taking the average. Since KSCM is a direct classification metric, and NNMF is an optimisation technique that requires more input data than factorisation ranks, we created a reference library that can satisfy both algorithms. The reference contains 10 randomly chosen pixels from a swatch, of which, the 8-neighbourhood pixels are extracted for each pixel chosen. For KSCM the 8-neighbourhood pixels are each classified. Similarly, for NNMF the 8-neighbourhood from each of the 10 pixels are

used as input data, obtaining the weights for each one. Therefore, sample sizes for each method are equal when calculating the accuracy.

5.2.4 Measuring Sensitivity to Variance of Spectra for KSCM and NNMF

In the previous sub-sections, we have seen three different approaches to classification of mixed paints. However, we did not test for changes in paint consistency, that is, changes in the mixture ratios. In this test, we prepared two reference datasets; one that only includes homogeneous paint mixtures at a mass-to-mass ratio of 50:50, and the other includes individual paints and homogeneous paint mixtures. The goal is to apply both KSCM and NNMF to the problem of classifying mixed paints of 40:60 and 60:40 against mixed paints at 50:50, to investigate which algorithm is more robust when presented with a spectral variability challenge.

5.2.5 Ground Truth of Paints

Like any other form of accuracy measurement, ground truth data is required. Since we created the mock-up samples shown in Figure 5.11 ourselves, we know the paints in each image, therefore, we can label them. The ground truth of paints for each image is illustrated in Figure 5.12 and Figure 5.13, showing layered paint images and homogeneously mixed paint images, respectively. Mixed regions in Images 1 and 2 are in the form of layers, hence the mixing ratio is unknown. On the other hand, mixtures in Images 3 and 4 are homogeneously mixed paints. In Image 3, the mixture is at 50:50, whereas in Image 4, the mixtures are either 60:40 or 40:60 as illustrated in the respective figure.

5.2.6 Hierarchical Paint Analysis - Evaluation Results and Discussion

Results are shown both qualitatively and quantitatively for the first stage of the Hierarchical Paint Analysis method. For the second stage, a quantitative result is given for each mock-up sample. Then, a discussion on the results obtained with regards to the identification of individual paints, and the identification of mixed paints is presented.

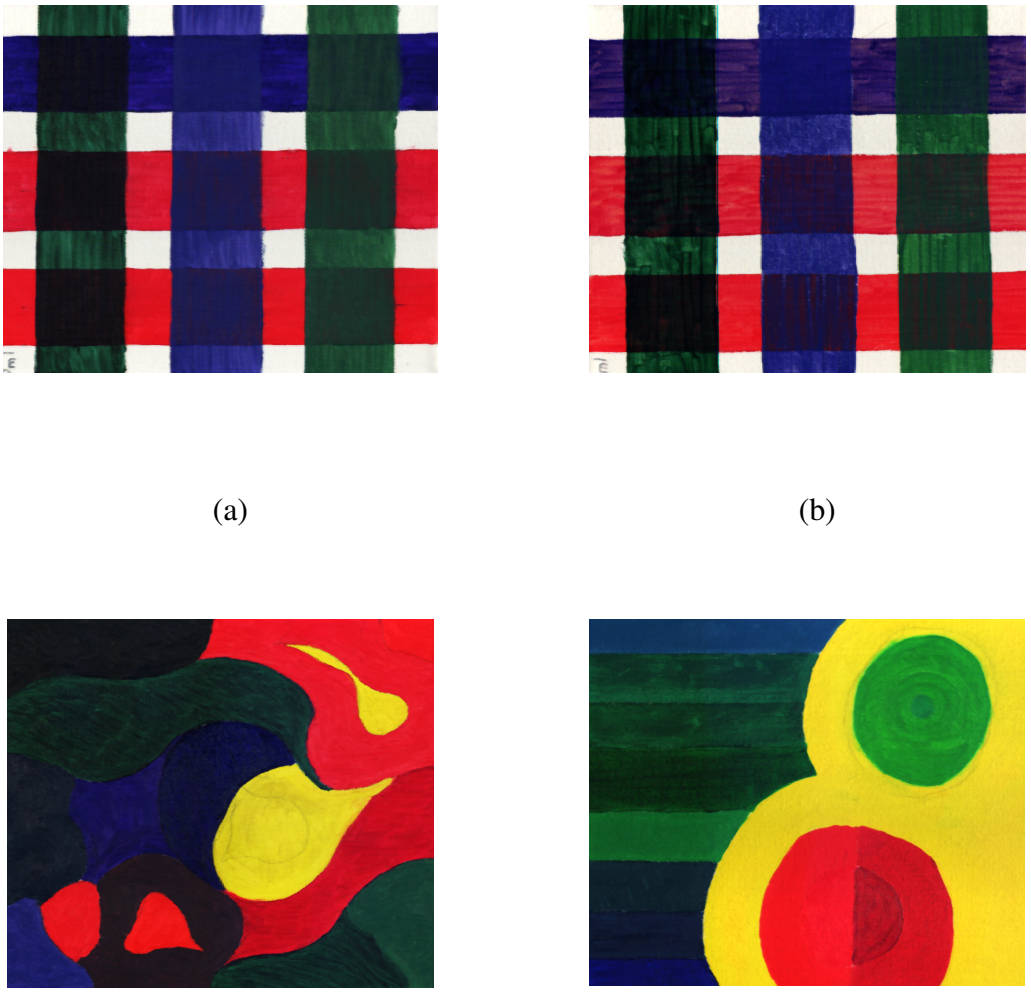

(a) (b)

(c) (d)

FIGURE 5.11: Paint RGB images (a) Image 1, (b) Image 2 (c) Image 3 and (d) Image 4.

5.2.6.1 KSCM Accuracy Results

Figure 5.14 illustrates paint region classification based on the SSM result output. The figure demonstrates that the metric that we propose, KSCM, is by far ahead of KSAM. As can be observed, KSCM has the highest true positive rate when compared to the other metrics. This means that KSCM has a much higher accuracy in correctly determining whether a region is an individual paint or a mixed paint. Table 5.3 sustains this claim by showing accuracy percentages for all mock-up samples. Similarly, Table 5.4 shows the accuracy for correctly identified paints in individual regions.

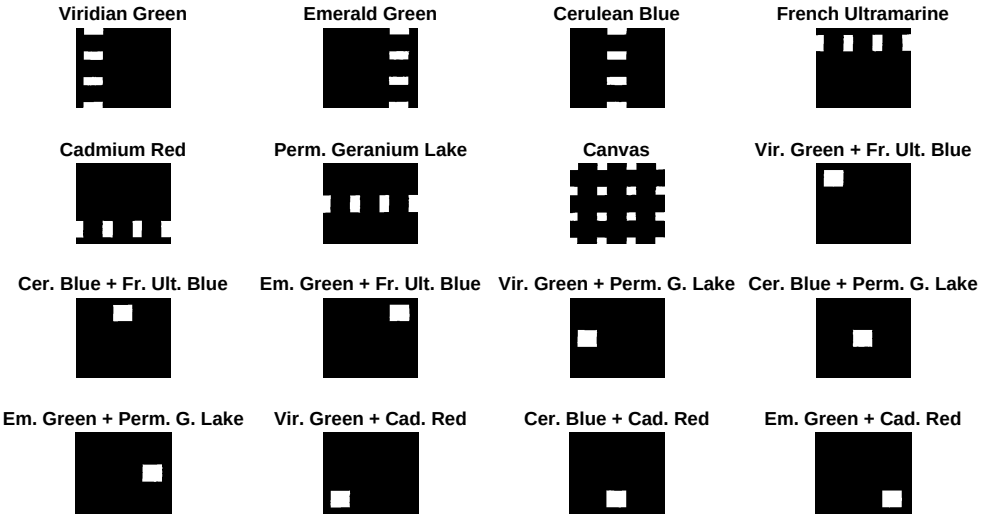

(a)

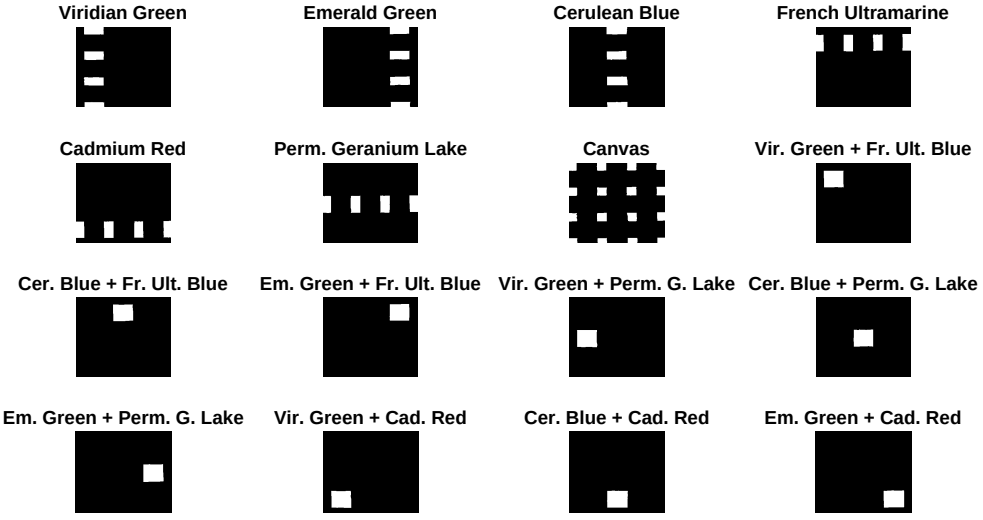

(b)

FIGURE 5.12: Paint ground truths for images containing layered paints (a) Image 1 and (b) Image 2.

Unfortunately, an application of KSAM for paint identification to provide a benchmark against which results can be compared could not be found. However, in [47], the authors compare their proposed KSAM with SAM, and report accuracy measures of 91.6% and 88.1%, respectively for their application in remote sensing. In results obtained in this work, the difference between the two methods is reflected, that is, the

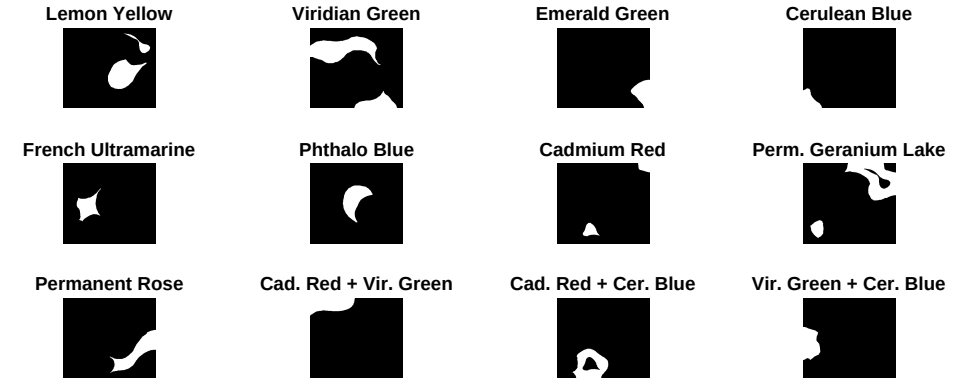

(a)

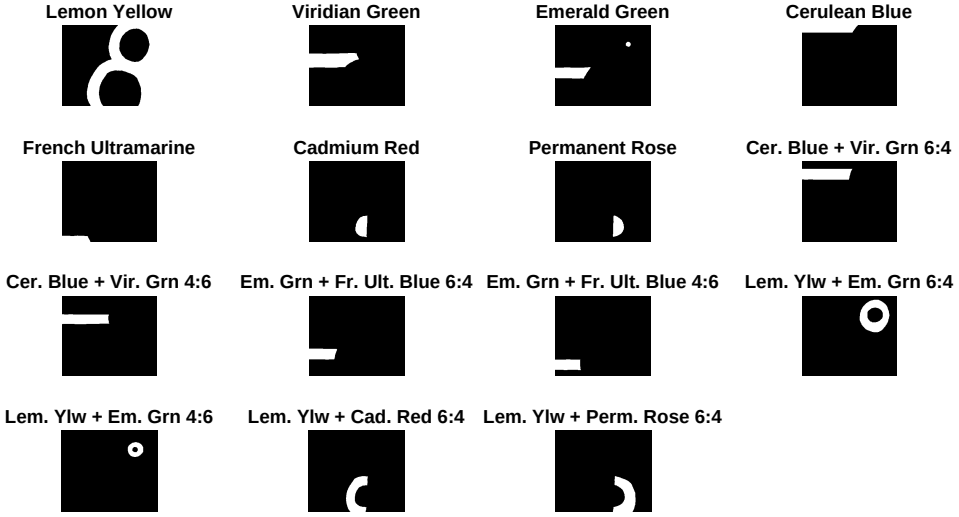

(b)

FIGURE 5.13: Paint ground truths for images containing homogeneously mixed paints (a) Image 3 and (b) Image 4.

KSAM algorithm is more accurate than SAM as shown in Table 5.3.

It is important to note that to determine mixed paint regions from individual paint regions, a threshold was set by qualitatively assessing the obtained results. The set thresholds were determined qualitatively for the ability to separate individual paints from mixed paints. These optimal thresholds for our application have been found to be; 0.3 for KSCM, 1.2 for KSAM, 0.95 for SCM, and 0.1 for SAM. Moreover, the same threshold is applied to each image sample.

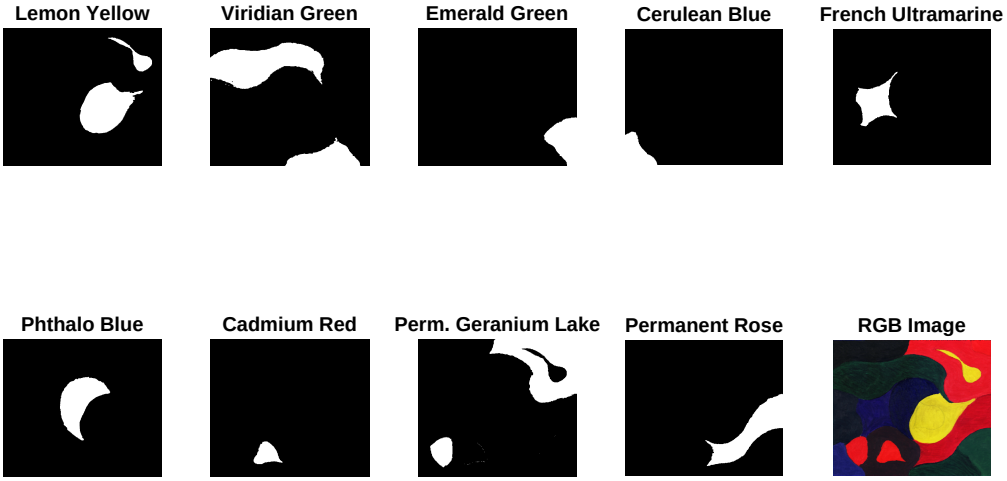

(a)

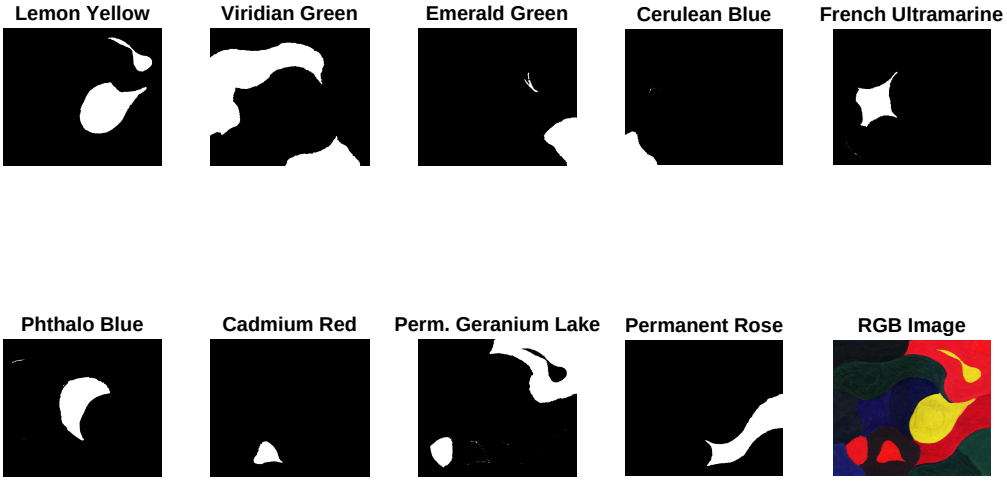

(b)

FIGURE 5.14: Classification of each feature in the form of binary labels, in which white patches refer to detected and black patches refer to not detected.The two metrics (a) KSCM and (b) KSAM are shown, demonstrating a higher accuracy for KSCM when compared to KSAM.

5.2.6.2 Hierarchical NNMF Accuracy Results

For the mixed paint identification performed in the hierarchical stage of this method, we show results for both class based classification and specific paint-based classification. Table 5.5 shows how a dataset of homogeneously mixed paint regions is classified by class and individual paint. Figure 5.15 presents an example of how the resulting detected mixture is selected to be the most similar to the input mixture. The figure shows a heatmap for the results obtained by the metric to determine which paint combination

TABLE 5.3: Comparing accuracy results for KSCM, KSAM, SCM and SAM for the problem of separating mixed regions from individual regions.

	Accuracy (%)			
	KSCM	KSAM	SCM	SAM
Image 1	95.9	93.5	89.8	81.6
Image 2	95.9	93.5	95.9	81.6
Image 3	100	93.7	93.7	87.5
Image 4	87.5	87.5	84.4	62.5

TABLE 5.4: Comparing accuracy results for KSCM, KSAM, SCM and SAM for the problem of correctly identifying the pure paints in individual regions.

	Accuracy (%)			
	KSCM	KSAM	SCM	SAM
Image 1	96.9	87.8	53.1	87.8
Image 2	93.9	86.7	54.1	89.8
Image 3	93.8	87.5	81.3	87.5
Image 4	90.6	43.8	43.8	43.8

correlates better with the input signature. Individual paints have been determined in a similar manner to Figure 5.15. Recall that the metric used for this identification includes both a correlation measure and an NNMF reconstruction error measure. Similarly, Table 5.6 show the results for paint mixtures in the form of layers. The reason why the two tables are separated is because light reacts differently when penetrating different paint layers, compared with homogeneously mixed paints. Therefore, the difference in accuracy is highlighted better by separating them.

Table 5.5 shows that the identification of a paint mixture by class, has a high accuracy, whereas performing identification of the same paint mixture by specific paint results in lower. One of the reasons why this happens is because of the high similarity between paints, such as different shades of reds. When any of these red paints is mixed with a blue paint, the outcome is a brown hue, in which the part of the spectral signature that represents the red will be shifted downwards or upwards in intensity, according to the signature of the blue paint. This causes a difficulty in determining which red hue was in the mixture. Since Image 3 contains only three paint mixtures, we included additional paint mixtures in the table, which may not be found in any of the showcased samples.

For the case of paint mixed by applying multiple layers - in our case two paints - we observe a considerable reduction in accuracy for both class identification and individual

TABLE 5.5: NNMF metric results for various homogeneous paint mixtures, separating them by class.

Ground Truth	Detected Classes	No. of Correctly Detected Classes	Detected Paints	No. of Correctly Detected Paints
LY + CRH	Yellow + Red	2	LY + PR	1
LY + VG	Yellow + Green	2	LY + VG	2
LY + CBH	Green + Red	0	VG + PGL	0
LY + EG	Yellow + Green	2	LY + VG	1
PB + LY	Green + Blue	1	VG + PB	1
PB + PR	Green + Blue	1	VG + FU	0
LY + PR	Yellow + Red	2	LY + PR	2
PR + EG	Red + Blue	1	CBH + PGL	0
FU + EG	Blue + Green	2	FU + EG	2
CRH + VG	Red + Green	2	PR + VG	1
CRH + CBH	Red + Blue	2	PGL + CBH	1
VG + CBH	Green + Blue	2	EG + CBH	1
	Accuracy	79.2%	Accuracy	50%

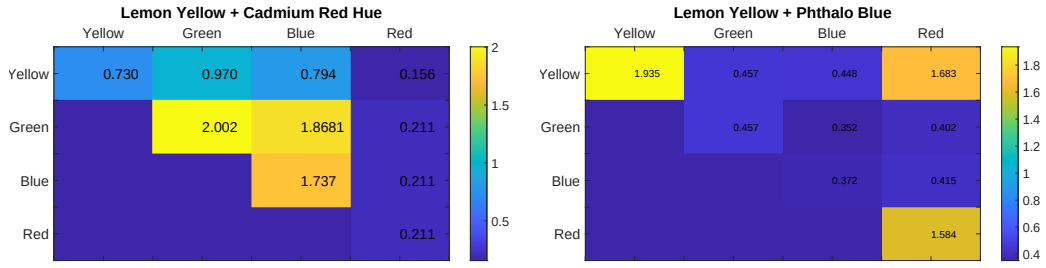

(a) (b)

FIGURE 5.15: Two examples of symmetrical matrices obtained from NNMF in the Hierarchical Paint Analysis method. The matrices represent (a) a result for a Yellow and Red mixture which correctly determines the Yellow and Red classes and (b) a result for a Yellow and Blue mixture which does not correctly determine the Yellow and Blue classes. Our metric represents the best result by the minimum value.

paint identification, as recorded in Table 5.6. However, we also observe that, for some paints, light penetration is stronger than for other paints. This may occur either due to the pigments' mixing ratio with the binder, or due to other physical properties of the pigments such as different sizes of particles. In fact, Figure 5.16 illustrates a KSCM classification result for Image 1, demonstrating that the layered patch comprised of Permanent Geranium Lake (bottom layer) and Emerald Green (top layer) is strongly correlated with Emerald Green as an individual paint.

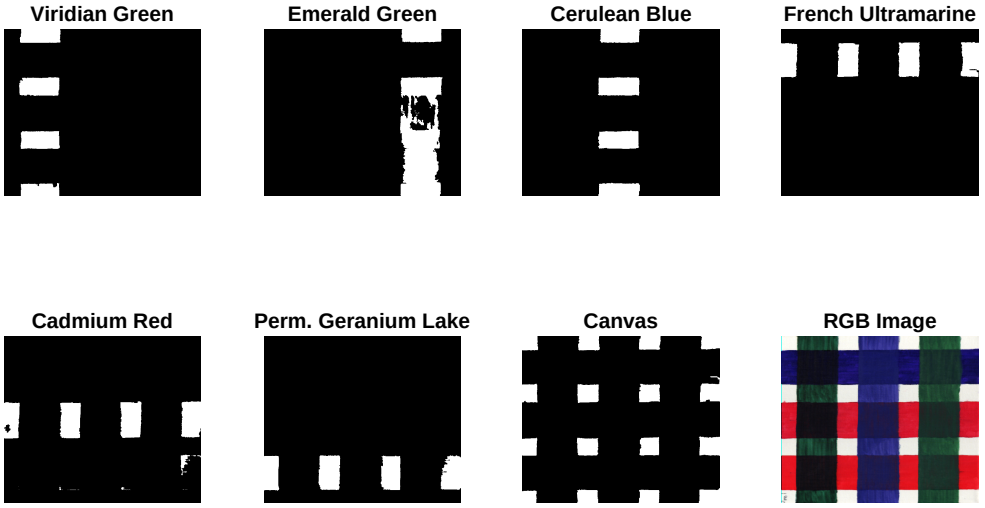

FIGURE 5.16: Demonstration to show that layered paints are harder to distinguish as mixtures when the top paint layer is highly reflective, where white indicates the presence of the paint/canvas while black indicates its absence.

TABLE 5.6: NNMF metric results for various layered paint mixtures, separating them by class

Ground Truth	Detected Classes	No. of Correctly Detected Classes	Detected Paints	No. of Correctly Detected Paints
VG + FU	Green + Blue	2	VG + FU	2
VG + CRH	Green + Blue	1	VG + FU	1
VG + PGL	Blue + Blue	0	PB + PB	0
FU + CBH	Green + Blue	1	EG + FU	1
FU + CRH	Blue + Yellow	1	LY + PB	0
FU + PGL	Blue + Green	1	FU + EG	1
EG + FU	Green + Yellow	1	LY + VG	0
EG + CRH	Green + Yellow	1	LY + VG	0
EG + PGL	Green + Yellow	1	LY + VG	0
	Accuracy	50%	Accuracy	27.8%

Having observed results for Hierarchical Paint Analysis, we can conclude that individual paint regions can be more accurately classified using KSCM rather than KSAM, SCM, or SAM. Consequently, KSCM is also the best metric to distinguish between mixed and individual regions. With regards to the second stage of the algorithm, the metric that

TABLE 5.7: Accuracy results for Global NNMF Paint Analysis method on both classification by class and by individual paint.

	NNMF Accuracy (%)	
	Class	Individual
Image 1	79.1	77.5
Image 2	80.0	75.3
Image 3	90.5	69.2
Image 4	95.5	82.1

combines a correlation and reconstruction error shows that homogeneous mixed paint signatures can be classified by class at an accuracy of 79.2%, and by individual paint at an accuracy of 50%. For layered mixed paints, an accuracy of 50% is achieved for separation by class, and an accuracy of 27.8% is achieved for separation by individual paint.

5.2.7 Global NNMF Paint Analysis - Quantitative Evaluation Results and Discussion

This method was inspired by the work in [20]. In fact, it follows a similar methodology except that, in our method, the KS coefficients are calculated externally such that we can use a linear model. In [20], the method was applied to classes of paints, therefore, it is unknown whether this method can be applied successfully to mixed paints, or individual paints of the same class. Hence, we apply the method to the problem of paint classification by individual paints of the same class, and also test for the capability of separating mixed paint regions by individual paints.

5.2.7.1 Global Optimisation NMF Accuracy Results

Two reference libraries are used for this method. The first library includes the signatures of each class, that is, yellow, blue, green, red and white. The second reference library includes each individual paint.

Figure 5.17 illustrates a visual representation for the separation by class of Image 3. The figure demonstrates that separation by class performs better for individual paint regions, rather than for mixed paint region separation. Table 5.7 includes accuracy measures for paint separation by class and individual paints, for the four images. The accuracy measures recorded in the table suggest that while individual paint accuracy is highly accurate, even for similar paints such as reds, the separation of paints in mixtures

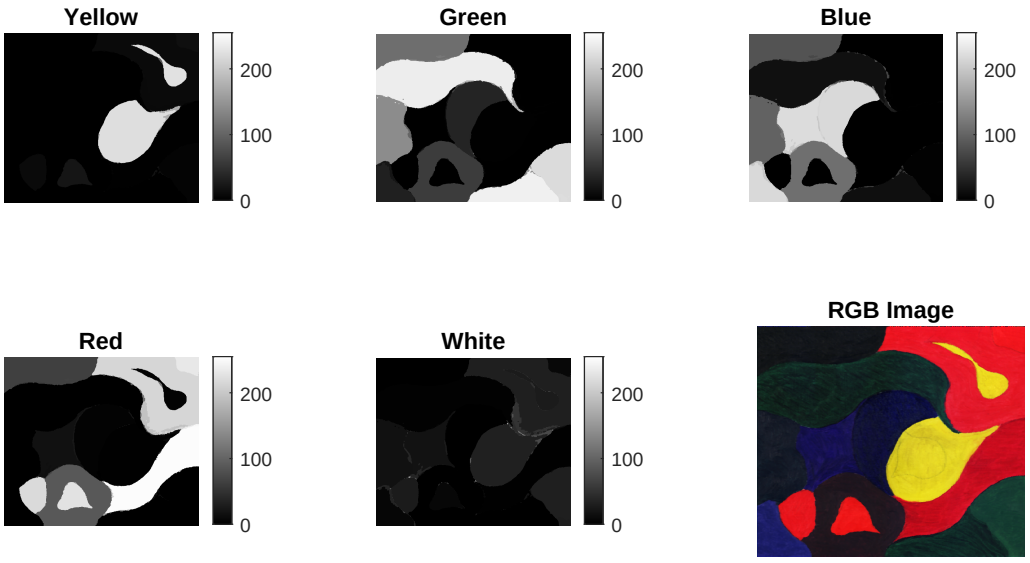

FIGURE 5.17: Classification result of Image 3 when separating the paint regions by paint class, where a brighter intensity indicates a higher class probability

is a difficult task, even for separation by class. Although our mixtures were mixed at a 50:50 mass-to-mass ratio, some of these mixtures are still dominated by one of the paints. This happens especially for mixtures that are mixed with a yellow paint. Other mixtures, for example those made up of Cadmium Red Hue and Viridian Green, appear as a very dark green mixture. These non-linear relationships of mixed paints cause the separation to be very difficult. Moreover, some of these paint combinations result in a paint signature that is very similar to other individual paints, such as the signature of the mixture made up of Cadmium Red Hue and Viridian Green, that is very similar to the signature the individual paint Phthalo Blue.

Similarly, Figure 5.18 illustrates the result obtained for Image 3, when the Global Optimisation NNMF is applied with the aim to separate the regions by individual paints. As can be observed, the separation for individual paint regions still performs better when compared to class identification of the same paint regions. Similarly, the separation of mixed regions is less accurate when identifying the mixed regions by the combination of individual paints.

Based on the achieved results, we can conclude that the Global Paint Region Analysis method accurately identifies individual paint regions, while identifying mixed paint

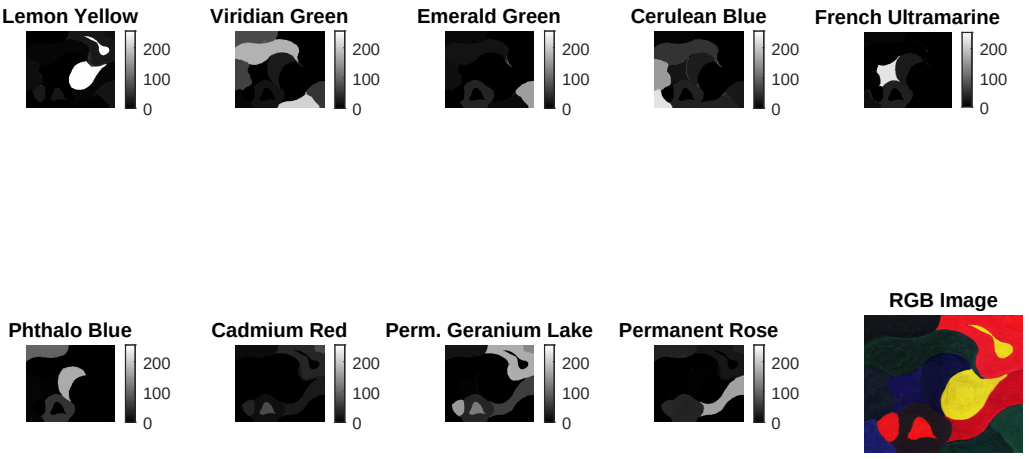

FIGURE 5.18: Classification result of Image 3 when separating the paint regions by individual paints.

regions with less accuracy. However, we notice that for some mixed regions, misclassifications occur because of the nature of the spectral signature, which strongly relates to one specific paint in the mixture as discussed in the previous section. Recalling that Images 1 and 2 include paint layer mixtures, and Images 3 and 4 include homogeneously mixed paint regions, we can notice an accuracy gap, caused by how light is interacting with the layered paints, causing weight biasing toward to the top paint layer. The lower accuracy observed in Image 4 for NNMF is due to the fact that the mixtures are not combined at a ratio of 50:50 but rather a variant of 60:40 or 40:60, as will be explained in more detail in Sub-Section 5.2.9.

5.2.8 Direct Classification of Mixtures - Quantitative Evaluation Results and Discussion

The Direct Classification of Paint Mixtures differs from the two methods mentioned previously. It still uses the same metrics, KSCM and NNMF separately, to measure the accuracy at classifying mixed paint regions. We first illustrate results obtained for classification of homogeneously mixed paint regions. Then we demonstrate the effect of spectral variance on the accuracy values of mixed paint classification.

TABLE 5.8: Accuracy results for Direct Classification of Paint Mixtures.

	Accuracy (%)	
	KSCM	NNMF
Image 1	89.7	71.4
Image 2	71.4	71.4
Image 3	87.5	81.3
Image 4	87.5	56.3

5.2.8.1 Mixed Paint Reference Library Results

The results obtained in this section demonstrate that it is possible to identify mixed paints by including a mixed paint reference in the library. Figure 5.19 illustrates the result obtained by KSCM after including the mixed paints in the library. In previous methods, we had nine individual paints, whereas now, we have nine individual paints and three types of homogeneously mixed paints. Therefore, our reference dataset will include 12 types of paint features. For this method, we do not include identification by classes as it would be difficult to separate mixed paints by class. Similarly, Figure 5.20 illustrates the obtained result for Image 3 by applying NNMF. We observe that the KSCM is better than NNMF in this case, because of the similarities between different mixed paints in terms of spectral properties, such as the amplitude of the signatures. On the other hand, due to the slight differences in spectral signature, KSCM is much more sensitive to these changes, therefore, yielding better results.

Table 5.8 demonstrates that the accuracy values achieved by both algorithms is high, with an average of 84% and 70.1% for KSCM and NNMF, respectively. As mentioned, when individual paints are mixed together, the resulting mixture sometimes appears to look similarly to a single individual paint. This can be accepted, because technically, as discussed in Chapter 4 in Section 2.2, individual paints are oftentimes created by mixing different pigments together. Essentially, by mixing individual paints, we are also mixing different pigments together to create different paints. Hence, this may explain the reduction in accuracy when comparing NNMF with KSCM.

5.2.9 Sensitivity to Variance for KSCM and NNMF Results

In this test, we observe that both KSCM and NNMF are affected by spectral variance, which is simulated by changing the mass-to-mass ratios of mixtures. Table 5.9 demonstrates that for each paint mixture, the accuracy decreases when the input paint mixture

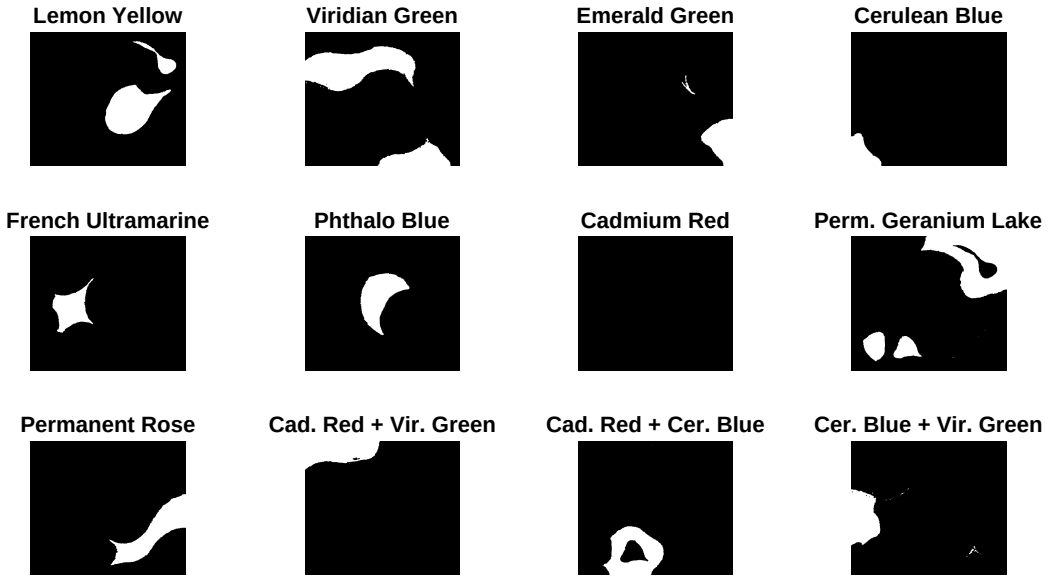

FIGURE 5.19: Classification result of Image 3 for KSCM. A mixed paint reference is provided in the reference library.

contains a variety of signatures of the same paint, including mixtures at 50:50, 60:40, and 40:60 mass-to-mass ratios.

For KSCM, the change in accuracy is much less than in NNMF. The reason behind this is that, for KSCM although there is a change in both absorbance and scattering coefficients, the spectral signature remains very similar. On the other hand, for NNMF although the spectral signature remains similar, the change in amplitude causes a higher error in the cost function. This lowers the accuracy in NNMF. This means that, although NNMF has the advantage of providing a weighted result, in which usually the maximum weight is taken to correspond to the best feature, it could also be that a small change in the spectra causes a misclassification. On the other hand, KSCM is minimally affected by these small changes because the main features between the two signatures are still strongly correlated.

5.2.10 Comparing the Classification Methods

In this section we have discussed three methods; the Hierarchical Paint Analysis approach, a Global NNMF Paint Analysis approach, and a Direct Classification of Mixtures

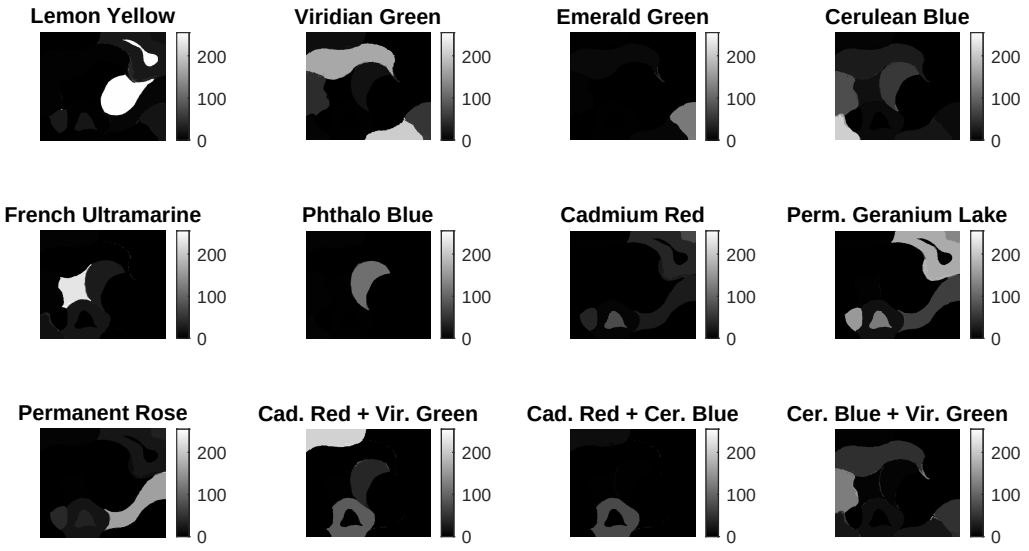

FIGURE 5.20: Classification result of Image 3 for NNMF. A mixed paint reference is provided in the reference library.

approach. Each method has its advantages, and in this sub-section we will be discussing these advantages, also considering the accuracy results that each method has achieved.

The Hierarchical Paint Region Analysis approach is very useful for applications in which no information is known regarding mixed paints. Its aim was to be able to find mixed paint regions, then separate those paints in the mixtures. The results show that the method performs well to separate the individual regions from the mixed regions, achieving an average accuracy of 93.3%, but suffers a loss in accuracy when identifying the individual paints in mixtures, when compared to the accuracy obtained by the Global NNMF Paint Analysis. In fact, the Hierarchical approach results in an accuracy of 79.1% for homogeneous mixtures while the Global NNMF approach results in an average accuracy of 93%. Similarly, accuracy measures for layered mixtures are 50% and 79.55% for the same method, respectively, when identifying by classes. Both algorithms follow a similar change in accuracy measures when identifying individual paints. The advantage of using the Hierarchical approach is that it returns results in the form of labels, without requiring any other form of interpretation. Meanwhile, the Global NNMF approach returns results in the form of weights, in which additional interpretation - visual or algorithmic - would be required to determine whether a paint region is a mixture of paints, and which paints are assigned to make up the mixture.

TABLE 5.9: Sensitivity to variance for KSCM and NNMF results for both KSCM and NNMF, when applied to a 12-Paint reference dataset that only includes mixed paints, and a 22-Paint reference dataset that also includes the individual paints.

Mixture	Accuracy (%)			
	NNMF		KSCM	
	12-Paint Ref	22-Paint Ref	12-Paint Ref	22-Paint Ref
LY + CRH	100	33.3	100	100
LY + VG	56.7	0	93.3	56.7
LY + CBH	0	100	100	100
LY + EG	100	100	100	100
PB + LY	100	100	83.3	83.3
PB + PR	13.3	23.3	20.0	20.0
LY + PR	96.7	53.3	100	100
PR + EG	46.7	60.0	30.0	30.0
FU + EG	20.0	0	73.3	63.3
CRH + VG	60.0	53.3	66.7	66.7
CRH + CBH	0	50.0	80	80.0
VG + CBH	100	0	100	100
Average	57.8	47.7	78.9	75.84

The third method, the Direct Classification of Mixtures, is an approach in which all information regarding the data is known. This means that, the individual paints used for the painting are known, and the types of mixtures, including their mixing ratios, are also known. This allows for a reference dataset that includes each and every paint in the image. Results for this method have shown that both KSCM and NNMF are viable candidates for this application. However, from our analysis, we have observed that when some paints are mixed, the resulting spectral signature is close in shape to the spectral signature of one of the paints. This specific observation mostly occured when using yellow as one of the paints in a mixture. For example, Figure 5.21 shows that when a yellow is mixed with a red, the spectral signature of the mixture remains similar to the signature of the red paint. On the other hand, when mixing a yellow paint with a blue paint, the result is a green paint. This caused the NNMF used for this application to fail in distinguishing between individual paints, such as Emerald Green, and mixed paints such as Cerulean Blue with Lemon Yellow. This problem also affects KSCM, but on a smaller degree.

Since the KS dataset obtained from [19] contained coefficients in the range of 400-750nm, we were not able to obtain these coefficients for the full range of our spectral camera, which was 400-1000nm. Therefore, we did have a loss of valuable information

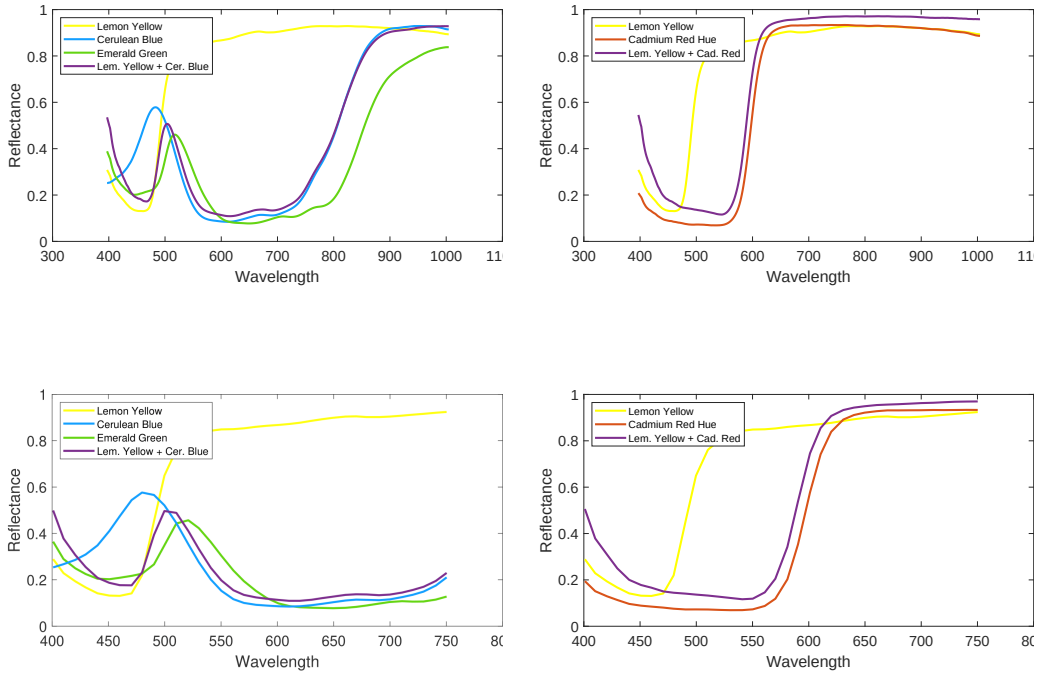

FIGURE 5.21: Similarity between a mixture of two paints and the individual paints used in the mixtures.

when deciding to use KS coefficients rather than R (reflectance) spectra. However, this suggests that if KS coefficients were to be simulated for the full range (400-1000nm), the results of mixture identification should improve further. This is because, as can be observed in Figure 5.21, the Lemon Yellow and Cerulean Blue mixtures correlate strongly with the original paints in the NIR region, rather than the non-existent Emerald Green paint, which shows high correlation with the mixture in the visible spectrum region.

Overall, KSCM performs better when the application is not data-driven, whereas an optimisation technique like NNMF is a better approach when no information is known regarding the mixed paints. As discussed in Section 5.2.9, KSCM is also much less sensitive to data variance, making it more suitable for cases in which the mixed paint is known, but the mass-to-mass ratio of the mixture is unknown. Additional visual results for each method are shown in Appendix B.

5.3 Case Study

The National Gallery of Art Museum, Washington, provided us with a hyperspectral illuminated manuscript 'The Birth of John the Baptist'. The hyperspectral data was

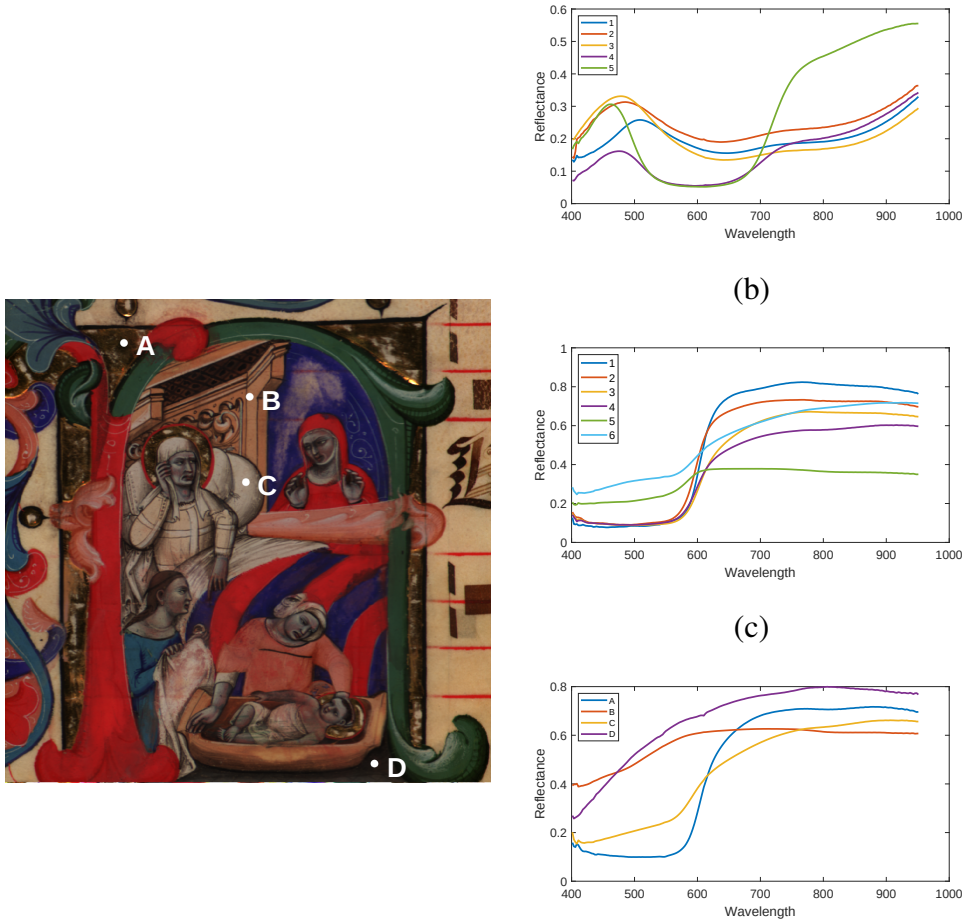

FIGURE 5.22: RGB representation of 'The Birth of John the Baptist'. White dots in (a) represent points of data used as additional reference points to be used for the classification application, with signatures are shown in (d). (b) Blue and Green signatures provided in dataset, and (c) Red and Pink signatures provided in dataset. The numbers in (b) and (c) match with the labels given to specific regions in [63].

captured in the range between 400-950nm, at a sampling rate of 2.5nm, resulting in a total of 220 wavelength channels. An RGB image of the illuminated manuscript is shown in Figure 5.22. As can be observed the manuscript contains multiple paints from different classes, such as reds, blues, greens and others. The manuscript includes many fine features which will prove to be a challenging task for our algorithms.

The data acquired consists of a cropped image derived from the original manuscript in [63], 11 paint reflectance signatures representing paints, such as blues, greens, reds and pinks, and their corresponding wavelengths. In [63], the results are shown qualitatively, and no data regarding the paint mappings, or identification accuracy was provided. Therefore, in this work, a qualitative analysis approach will be performed for both the segmentation and classification results.

FIGURE 5.23: A snapshot example of a fine detailed region taken from Figure 5.22.

In this study, the proposed SSM algorithm is applied to the case study with the goal of achieving a strong spectral segmentation. Since the data capture for this image was done at an unknown site with unknown ambient conditions, some changes to the threshold parameters were needed because initial conditions during image acquisition may vary from those in our samples.

5.3.1 SSM - Case Study

Since the image being studied in Figure 5.22 includes many fine features, it was challenging to decide which initial size of superpixels is ideal. If the line-like features are important and must be captured, then the initial size of superpixels must be small enough to satisfy the Nyquist criterion for their size. If we zoom in, we can observe that each of these line-like features is about four pixels wide, as shown in Figure 5.23, meaning that each superpixel should be smaller than 2×2 pixels. This means that for SSM to be able to perform better, a higher resolution is required, such that a better representation of these fine features can be obtained.

The results for both SLIC and SSM are illustrated in Figure 5.24. The results show that SSM clearly merges regions of similar hues, but fails to process some of the other regions. This is because the threshold for the intra-class variance is selected manually, and reducing it even further may affect currently stable regions. On the other hand, the algorithm provides a clean segmentation result, highlighting areas of similar hue. In the next couple of paragraphs, an analysis and discussion on a fine feature region and a region containing similar adjacent hues region is given.

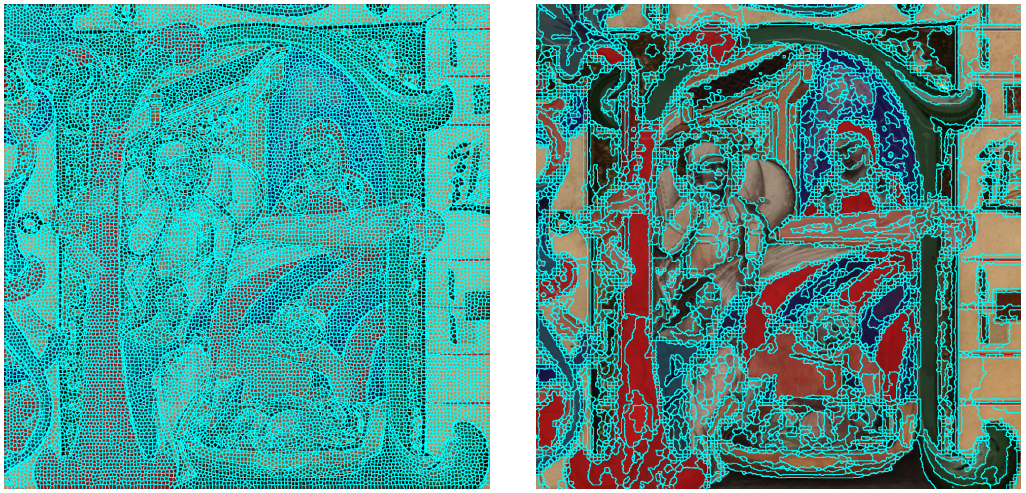

(a) (b)

FIGURE 5.24: Segmentation results for (a)SLIC and (b) SSM for case study image.

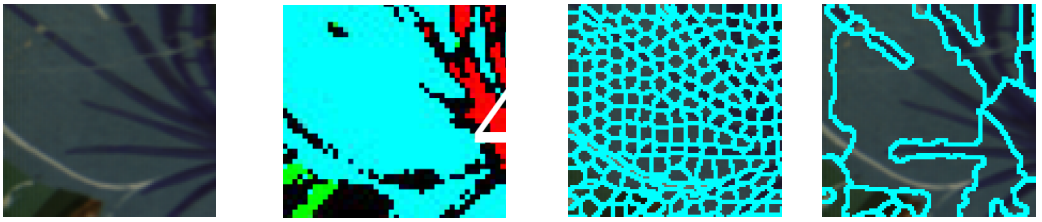

(a) RGB Image (b) Classification [63] (c) SLIC (d) SSM

FIGURE 5.25: Results obtained for a fine detail region (a). Image (b) is a cropped map image reproduced from the blues and greens map in [63], (c) is the SLIC result and (d) is the SSM result.

Figure 5.25 illustrates the result for the fine-detailed region, for both SSM and SLIC. At first glance, the SLIC result appears to be much cleaner than SSM, at the cost of over-segmentation showing clearly the areas containing fine details. On the other hand, SSM shows these regions less clearly. We suspect that this happens because the algorithm is heavily focused on correlation, and when these initial superpixels are very small, artifacts between neighbouring pixels may affect correlation, especially when not enough pixels are being used to create a representative signature. However, the results obtained by SSM closely match the results obtained in [63] for this particular region, whereas in SLIC this is unclear. In both cases presented in Figure 5.24, the initial number of superpixels for both algorithms should be 260,500 superpixels to satisfy the Nyquist criterion. However, this is very computationally expensive. Therefore, the total number of superpixels was set to 20,000 instead. Moreover, as mentioned earlier in this section, SSM should be

FIGURE 5.26: A snapshot example of a region painted with similar paints taken from Figure 5.22.

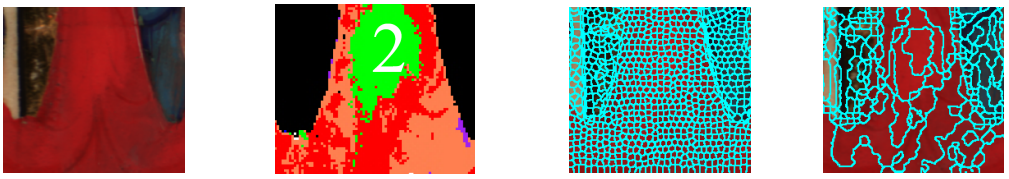

(a) RGB Image (b) Classification [63] (c) SLIC (d) SSM

FIGURE 5.27: Results obtained for a red region containing multiple red paints (a). Image (b) is a cropped map image reproduced from the reds and pinks map in [63], (c) is the SLIC result and (d) is the SSM result.

able to perform much better given that a higher resolution image for these fine features is used.

Figure 5.27 illustrates the result of both SSM and SLIC for the red region which does not contain fine details, shown in Figure 5.26. As can be observed, SSM in this case manages to capture the boundary between regions clearly and reduces over-segmentation when compared to SLIC. When the results of SSM and SLIC are compared to the classification of paints obtained by [63] for 'The Birth of John the Baptist', we can observe that SSM manages to capture the regions made up of different paints, whereas for SLIC this difference is unclear.

From these observations, we conclude that our algorithm is superior in terms of providing a much clearer representation of the paints in the image by reducing over-segmentation, and also, in its capability of capturing different but very similar paints in adjacent regions. However, this comes at the compromise of requiring a high resolution image for an optimal performance on fine features.

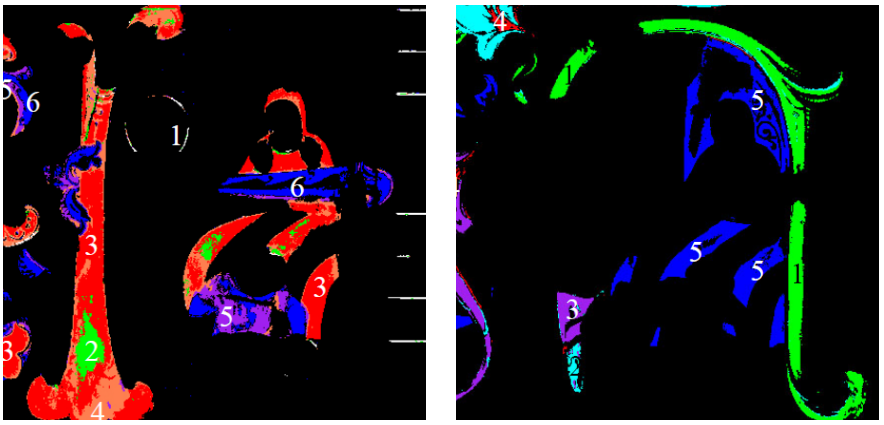

(a) Reds Classification [63] (b) Blues Classification [63]

FIGURE 5.28: Classification results obtained from [63]. (a) Represents reds visual ground truth (b) Represents blues visual ground truth.

5.3.2 KSCM vs Global NNMF - Case Study

The painting 'The Birth of John the Baptist' contains a variety of different paints [63], in which the authors identify a selection of blues, greens, reds and pink paints in a pixel-by-pixel manner. The results obtained in [63], illustrated in Figure 5.28, are used as a visual ground truth of the paint locations. These results were not provided as data, but only as images and, hence, they could only be used for visual comparison. To present an accurate comparison, matching labels were assigned, such that a Red 1 in [63] matches with a Red 1 map result obtained from this analysis.

As can be observed, although not too clearly, there are various regions in which paints seem to be overlapping, as also stated by the authors of [63]. Whether the painted regions in 'The Birth of John the Baptist' are a result of homogeneous or layered mixtures of paints makes no difference at this stage, since the ground truth spectra that we were provided with have been obtained directly from the painting [63]. Therefore, the paint signatures are treated as different individual paints, as done in the Direct Classification of Mixtures method. Hence, mixed regions cannot be determined with our current data. As done for the Direct Classification of Mixtures, we apply both KSCM and NNMF to this problem.

Figure 5.29 illustrates the results obtained by KSCM showing that the accuracy is very high when compared to the ground truth. One can observe how, although many of the different paints are similar, KSCM is still capable of classifying the different regions correctly. The thresholds used for reds is 0.55 and for blues, is 0.5. When compared to the result of [63], it is observed that Reds 5 and 6 are captured as one paint in our classification, and 5 is detected at completely unrelated areas. After a deeper analysis,

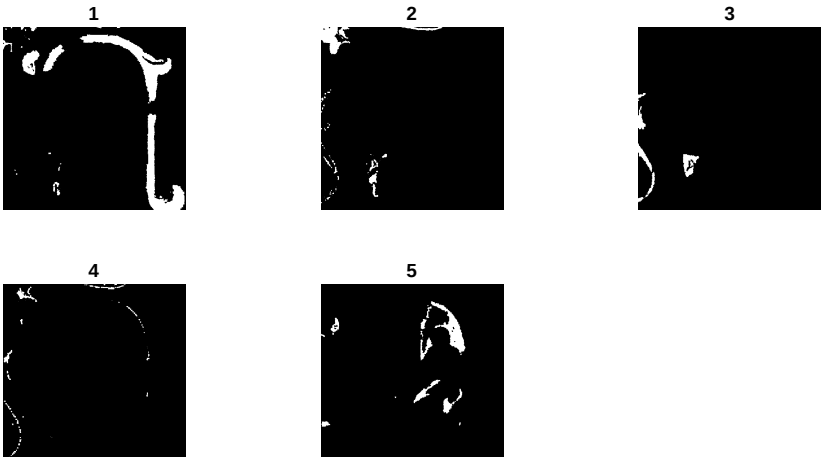

(a) Blues

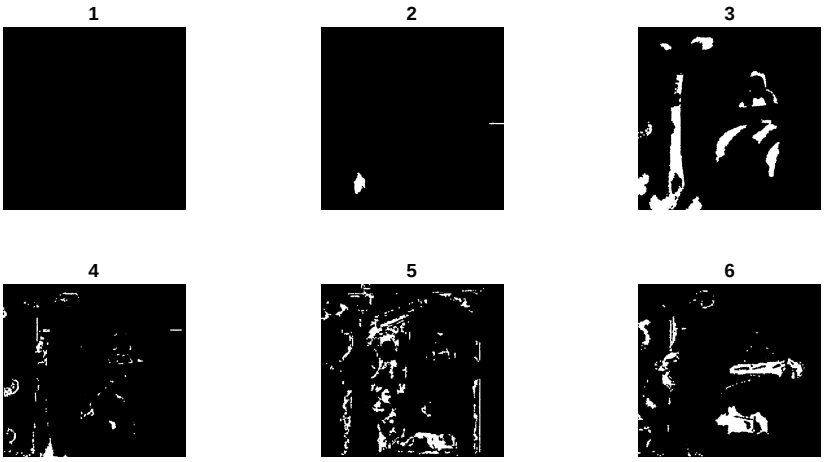

(b) Reds

FIGURE 5.29: KSCM classification results for (a) blue paints and (b) red paints. The numbered features correspond with the signatures in Figure 5.22

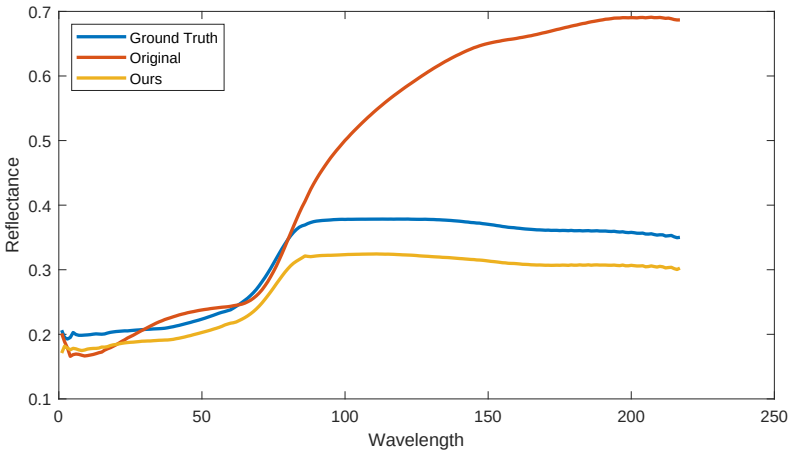

FIGURE 5.30: Illustrating the difference between the red paint reference number 5, and the actual red paint obtained from a pixel in the region mapped as 5, and comparing it to a paint obtained from a pixel in the region of John the Baptist's skin.

it was observed that the red paint reference 5 does not match with a pixel signature obtained from the region mapped as 5 in [63]. Figure 5.30 illustrates that the red paint reference 5 is more similar to a region where John the Baptist is lying. This means that the result obtained by KSCM was not incorrect for this area.

For the case of NNMF, the implementation is performed using blue (includes green) and red paints (includes pinks) obtained from [19]. Other paint references were added: a base (paper), a dark brown paint, a light brown paint, and a white paint, collected from the painting locations shown in Figure 5.22. Figure 5.31 illustrates the obtained results, focusing only on the blue (includes green), red (includes pink), and other features. As observed, NNMF did not manage to properly separate the different paints because it is much more sensitive to similar spectra.

As discussed in Subsection 5.2.9, the KSCM is less sensitive to spectral variance allowing for a more accurate classification of these different paint regions having similar spectra. Therefore, we conclude that KSCM is much better than NNMF when classifying paints having very similar spectra.

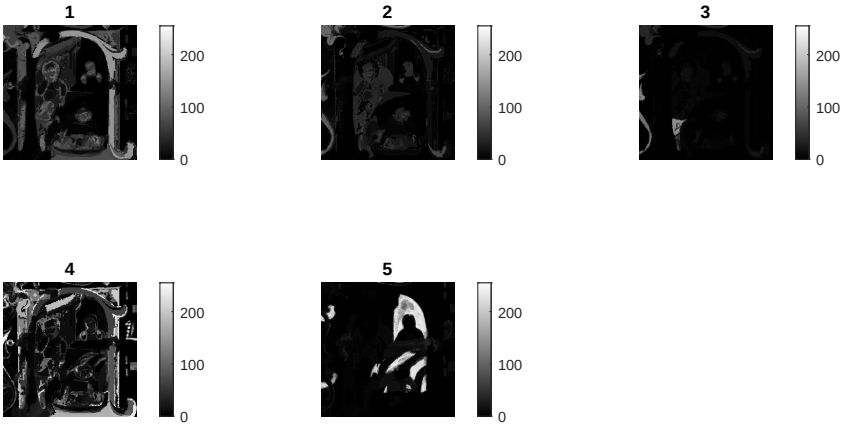

(a) Blues

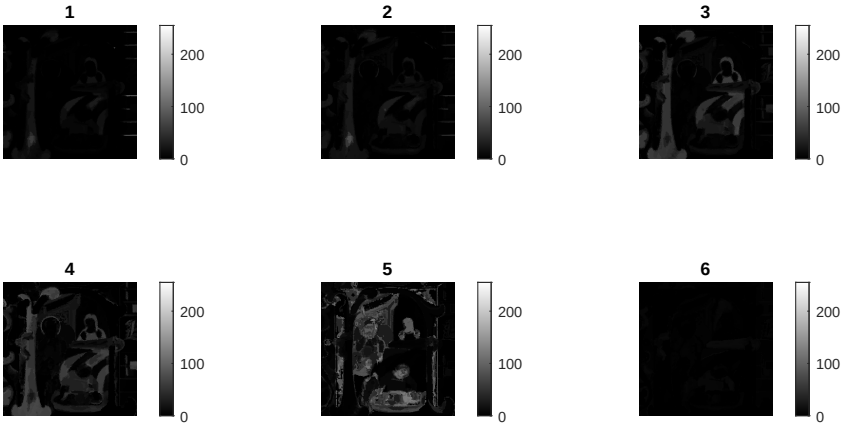

(b) Reds

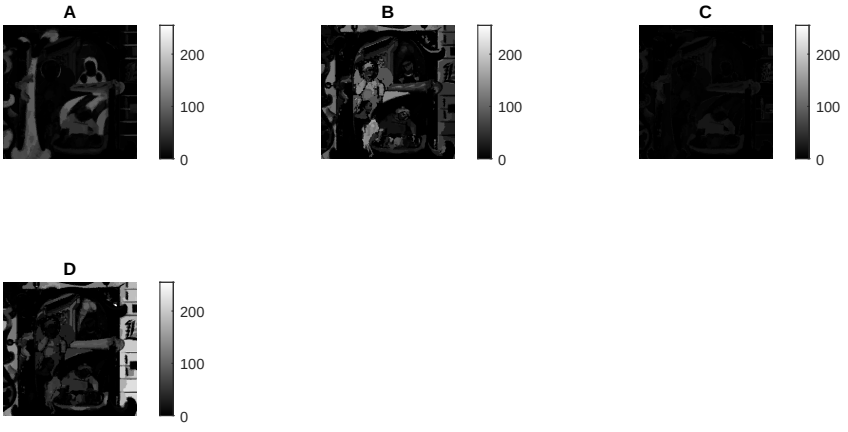

(c) Others

FIGURE 5.31: NNMF Classification results for (a) blue paints (b) red paints and (c) Other paints. The numbered features correspond with the signatures in Figure 5.22

Chapter 6

Conclusion

The restoration of paintings is a very important subject because it involves saving both privately owned heritages and historical artworks related to culture. To perform restoration of paintings a sequence of analytical steps are required taking into consideration every possible material on the painting, and running tests to determine what these materials consist of. There are two main approaches one can take to perform data acquisition regarding the sample. The first approach is by using invasive techniques which require taking samples directly from the painting. This approach is time consuming and requires extra care when handling. Non-invasive methods are a secondary approach to taking sample data of a painting. These approaches are typically based on spectroscopy techniques which measure energy being radiated from the surface of the painting. The non-invasive approach to paint identification for restoration of artworks is an ongoing area of research, and applied in this dissertation.

Spectroscopy is the study of the relationship between matter and electromagnetic waves. In this dissertation, the typical range of wavelengths that we study range between the visible spectrum (400nm) and the near infra-red region (until 1000nm). There are various spectroscopy techniques that one can study to solve the problem of paint identification. Hyperspectral Imaging (HSI) is one type of spectroscopy techniques in which, depending on the resolution, hundreds of reflection measurements are observed within the spectral range, providing a smooth and well defined spectral signature of a paint.

In this work, a pipeline to extract features and classify paints on a painting is presented. This pipeline was developed with the aim to aid restorers analyse paintings with ease, without having to spend an excessive amount of time trying to interpret the data.

This dissertation presents two main contributions. The first contribution is that a segmentation approach is presented which focuses on separating paint regions based

on the homogeneity of spectra. Secondly, three mixed paint separation techniques are presented, evaluated and compared. The segmentation part of this work was published in the ICIP conference, 2021 [64].

Hyperspectral images of the mock up samples were acquired and pre-processed. Paint similarities and differences were studied in which it was found that KSAM separated individual paints much better than other traditional spectral mappers such as SAM and SCM. However, KSAM under performed when different thickness regions of the same paints were being separated. To mitigate this under-performance we modified the KSAM, by applying correlation to the kernel rather than the spectral angle resulting in KSCM. The proposed segmentation approach consists of three stages, each stage contributing to a lower over-segmentation result while increasing the homogeneity of the regions by making use of KSCM.

Three methods were developed for the separation of paints in mixed regions. The initial analysis was performed on homogeneously mixed regions, as not enough information regarding layered regions, such as layer thickness, was available. The three methods follow a different approach, with the aim to perform a comparative analysis and determine which approach is best suited for the separation of mixed paints.

The results for the proposed segmentation algorithm are qualitatively and quantitatively compared to segmentation results obtained for SLIC. Results obtained show that the proposed algorithm is capable of separating similar hue paint regions with an average F-Score of 48.5% while keeping the over-segmentation at an average of 67.3%. When compared to SLIC, these results mean an average improvement of 39.8% in F-Score and an average improvement of 27.5% in over-segmentation. Qualitative observations show that the proposed algorithm can provide a much clearer image for the restorer than the SLIC algorithm. As also observed, over-segmentation is not eliminated completely, and a very small amount of under-segmentation may be obtained, in fact, for the proposed segmentation algorithm the under-segmentation increased by an average of 0.1% when compared to SLIC. This means that although the proposed algorithm performs better than state-of-the-art algorithms such as SLIC, improvements to over-segmentation can still be made following upon this work.

For the classification stage, the three algorithms showed that it is very possible to separate individual paints from mixed paints signatures. Two methods consisted of separating individual paints from mixtures by optimising for all possibilities of the initial paint palette. The third method considers mixed paints to be individual paints, and classification is performed by directly referring to a database that contains both individual and mixed paint signatures. The first two methods show that separation of individual

paints from mixtures is possible. The first method, Hierarchical Paint Analysis, shows that by using metrics such as KSCM, the individual paint regions can first be separated from mixed paint regions. Then, an optimisation approach such as a local NNMF is used to determine which paints are present in each mixed region in a combination of features approach. Results for this method show an accuracy of 79.2% for separation by paint class, and 50% for separation by individual paints. The second method, Global Optimisation NNMF, employs a global optimisation NNMF approach, in which all individual paint features on the painting are used as features, and each region is weighted against these features. Results for this stage show that separation by class is achieved at an accuracy average of 86.3%, while separation by individual paints is achieved at an accuracy average of 76%. Finally, the third method, Direct Classification of Mixtures, shows promising results when using KSCM as the direct classifier, with an accuracy average of 84%, whereas the NNMF approach in this case achieved an accuracy average of 70%. Moreover, it was observed in experimentation discussed, that KSCM is less susceptible to spectral variance of these mixtures. In fact, KSCM managed to obtain an accuracy average of 75.8% when attempting to classify mixed paints directly with a varying mixing ratio that ranged between 40:60 to 60:40, with the reference being 50:50. Out of the three algorithms, The Hierarchical Paint Analysis method results in the poorest performance. The Global NNMF Paint Analysis had a higher accuracy for determining paint classes when compared to the Hierarchical Paint Analysis. The Direct Classification of Mixtures, presented the best performance, however, it works in a different manner as the mixtures are included in the reference library.

The proposed segmentation approach was applied to a case study 'The Birth of John the Baptist', in which the results are reflective to what was achieved in the mock-up samples, when compared to SLIC. Similarly, with regards to classification, the Direct Classification of Mixtures approach was used, and the results obtained for KSCM when compared to NNMF are also reflective of results obtained on our mock-up samples.

6.1 Improvements and Further Work

Overall, the performance of the proposed segmentation algorithm is considered to be satisfactory since the original goals that were set prior to design are well achieved. However, as discussed in Chapter 5, better performance was achieved for one of the mock-up samples in which very similar hues were painted adjacent to each other. The result that was obtained suggest that an adaptive or dynamic threshold can improve the results for paintings in which no over-segmentation of very distinct regions is desired

while managing to separate regions of highly similar hues. Therefore, this may be interesting to study since paint difference can vary throughout the image. Moreover, a higher resolution may also improve merging quality especially for fine features.

The range of wavelengths used for the classification methods was originally between 400nm and 1000nm. However, since experimentation suggested that the KS coefficients would result in a better performance, a narrower wavelength range was used, such that it would match with the database used for the KS coefficients. It was observed, that due to the shortening of the wavelength range, now between 400nm and 700nm, some paint features were being lost. Therefore, as a further improvement on this work, the full range, or a much wider range (Visible to Short-wave Infrared) should be considered, as this would also open analysis possibilities to the identification of binders in paints.

Currently, manual operation is required to operate the classification methods on the segmentation result. As a further improvement of the proposed methods, a tool could be implemented which would automate the entire process and enhance the quality of analysis which would further facilitate the analysis of paintings.

References

[1] Vitorino, T., Casini, A., Cucci, C., Melo, M., Picollo, M., and Stefani, L. Non-invasive identification of traditional red lake pigments in fourteenth to sixteenth centuries paintings through the use of hyperspectral imaging technique. *Applied Physics A*, 121, 07 2015. doi: 10.1007/s00339-015-9360-4.

[2] Cosentino, A. Identification of pigments by multispectral imaging: A flowchart method. *Heritage Science*, 2, 04 2014. doi: 10.1186/2050-7445-2-8.

[3] Attas, M., Cloutis, E., Collins, C., Goltz, D., Majzels, C., Mansfield, J., and Mantsch, H. Near-infrared spectroscopic imaging in art conservation: Investigation of drawing constituents. *Journal of Cultural Heritage*, 4:127–136, 04 2003. doi: 10.1016/S1296-2074(03)00024-4.

[4] Li, J. and Wan, X. Superpixel segmentation and pigment identification of colored relics based on visible spectral image. *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy*, 189, 08 2017. doi: 10.1016/j.saa.2017.08.042.

[5] Deborah, H., George, S., and Hardeberg, J. Pigment Mapping of the Scream (1893) Based on Hyperspectral Imaging. volume 8509, pages 247–256, 06 2014. ISBN 9783319079974. doi: 10.1007/978-3-319-07998-1_28.

[6] Vorobyev, M. Ecology and evolution of primate colour vision. *Clinical & experimental optometry : journal of the Australian Optometrical Association*, 87:230–8, 08 2004. doi: 10.1111/j.1444-0938.2004.tb05053.x.

[7] Academy, K. Spectroscopy: Interaction of light and matter. Website, 2015. URL https://bit.ly/3xperS0. Accessed May 2020.

[8] Daveri, A., Paziani, S., Marmion, M., Harju, H., Vidman, A., Azzarelli, M., and Vagnini, M. New perspectives in the non-invasive, in situ identification of painting materials: The advanced MWIR hyperspectral imaging. *Trends in Analytical Chemistry*, 98:143–148, 2018.

[9] Francisca, P., Meza, P., Pezoa, J. E., and Torres, S. Thermal Characterization of a NIR Hyperspectral Camera. pages 81–86, 10 2011. doi: 10.1117/12.898126.

[10] Born, M., Wolf, E., Bhatia, A. B., Clemmow, P. C., Gabor, D., Stokes, A. R., Taylor, A. M., Wayman, P. A., and Wilcock, W. L. *Basic properties of the electromagnetic field*, page 1–74. Cambridge University Press, 7 edition, 1999. doi: 10.1017/CBO9781139644181.010.

[11] Gueli, A., Bonfiglio, G., Pasquale, S., and Troja, S. Effect of particle size on pigments colour. *Color Research & Application*, 42, 06 2016. doi: 10.1002/col.22062.

[12] Sands, S. Pigment volume concentration and its role in color. Website, December 2016. URL https://bit.ly/3Aq1B85. Accessed May 2020.

[13] Delaney, J., Ricciardi, P., Glinsman, L., Facini, M., Thoury, M., Palmer, M., and de la Rie, R. Use of Imaging Spectroscopy, Fiber Optic Reflectance Spectroscopy and X-ray Fluorescence to Map and Identify Pigments in Illuminated Manuscripts. *Studies in Conservation*, 59, 01 2013. doi: 10.1179/2047058412Y.0000000078.

[14] Miljkovic, M., Djordjevic, D., Djokic, M., Miljkovic, V., and Stoilkovski, K. Qualitative analysis of pigments and binders in paint colors used in 'gvozden' by mica popovic. *Advanced technologies*, 6:56–59, 01 2017. doi: 10.5937/savteh1702056M.

[15] Ropret, P., Zoubek, R., Škapin, A., and Bukovec, P. Effects of ageing on different binders for retouching and on some binder–pigment combinations used for restoration of wall paintings. *Materials Characterization*, 58:1148–1159, 11 2007. doi: 10.1016/j.matchar.2007.04.027.

[16] Forsyth, D. A. and Ponce, J. *Computer Vision A Modern Approach*, chapter Light and Shading, pages 62–97. Pearson Education Limited, 2012.

[17] Berns, R. and Mohammadi, M. Evaluating Single- and Two-Constant Kubelka-Munk Turbid Media Theory for Instrumental-Based Inpainting. *Studies in Conservation*, 52:299–314, 12 2007. doi: 10.1179/sic.2007.52.4.299.

[18] Shuqiang, L., Liu, Y., Hou, M., Yin, Q., Wu, W., and Yang, X. Quantitative analysis of mixed pigments for Chinese paintings using the improved method of ratio spectra derivative spectrophotometry based on mode. *Heritage Science*, 8, 03 2020. doi: 10.1186/s40494-020-00372-5.

[19] Okumura, Y. Developing a Spectral and Colorimetric Database of Artist Paint Materials. Masters Thesis, Chester F. Carlson Center for Color Science of the College of Science Rochester Institute of Technology, 2005.

[20] Tan, J., DiVerdi, S., Lu, J., and Gingold, Y. I. Pigmento: Pigment-Based Image Analysis and Editing. *IEEE Transactions on Visualization and Computer Graphics*, 25:2791–2803, 2019.

[21] Daniel, F., Mounier, A., Pérez-Arantegui, J., Pardos, C., Prieto-Taboada, N., Vallejuelo, S., and Castro, K. Hyperspectral imaging applied to the analysis of Goya paintings in the Museum of Zaragoza (Spain). *Microchemical Journal*, 126:113–120, 12 2015. doi: 10.1016/j.microc.2015.11.044.

[22] Barni, M., Pelagotti, A., and Piva, A. Image processing for the analysis and conservation of paintings: Opportunities and challenges. *Signal Processing Magazine, IEEE*, 22:141 – 144, 10 2005. doi: 10.1109/MSP.2005.1511835.

[23] Brouwer, P. *Theory of XRF getting acquanted with the principles*, chapter Basics to XRF, pages 10–20. PANalytical B.V, 2010.

[24] Appolonia, L., Vaudan, D., Chatel, V., Aceto, M., and Mirti, P. Combined use of FORS, XRF and Raman spectroscopy in the study of mural paintings in the Aosta Valley (Italy). *Analytical and bioanalytical chemistry*, 395:2005–13, 09 2009. doi: 10.1007/s00216-009-3014-3.

[25] Elmasry, G. and Sun, D.-W. *CHAPTER 1. Principles of Hyperspectral Imaging Technology*, pages 3–43. 12 2010. ISBN 9780123747532. doi: 10.1016/B978-0-12-374753-2.10001-2.

[26] Cavaleri, T., Giovagnoli, A., and Nervo, M. Pigments and Mixtures Identification by Visible Reflectance Spectroscopy. *Procedia Chemistry*, 8:45–54, 12 2013. doi: 10.1016/j.proche.2013.03.007.

[27] Yao, H. and Lewis, D. *Spectral Preprocessing and Calibration Techniques*, pages 45–78. 12 2010. ISBN 9780123747532. doi: 10.1016/B978-0-12-374753-2.10002-4.

[28] Grabowski, B., Masarczyk, W., Głomb, P., and Mendys-Frodyma, A. Automatic pigment identification from hyperspectral data. *Journal of Cultural Heritage*, 31, 02 2018. doi: 10.1016/j.culher.2018.01.003.

[29] Wu, T., Li, G., Yang, Z., Zhang, H., Lei, Y., Wang, N., and Zhang, L. Shortwave Infrared Imaging Spectroscopy for Analysis of Ancient Paintings. *Applied Spectroscopy*, 71, 11 2016. doi: 10.1177/0003702816660724.

[30] Liu, Y., Shuqiang, L., Hou, M., and Yin, Q. The comparison between NMF and ICA in pigment mixture identification of ancient Chinese paintings. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-3:1169–1172, 04 2018. doi: 10.5194/isprs-archives-XLII-3-1169-2018.

[31] Rohani, N., Pouyet, E., Cossairt, O., Katsaggelos, A., and Walton, M. Nonlinear Unmixing of Hyperspectral Datasets for the Study of Painted Works of Art. *Angewandte Chemie*, 57, 06 2018. doi: 10.1002/anie.201805135.

[32] Kleynhans, T., Patterson, C., Dooley, K., Messinger, D., and Delaney, J. An Alternative Approach to Mapping Pigments in Paintings With Hyperspectral Reflectance Image Cubes Using Artificial Intelligence. *Heritage Science*, 07 2020. doi: 10.21203/rs.3.rs-40103/v1.

[33] Polak, A., Kelman, T., Murray, P., Marshall, S., Stothard, D., Eastaugh, N., and Eastaugh, F. Hyperspectral imaging combined with data classification techniques as an aid for artwork authentication. *Journal of Cultural Heritage*, 03 2017. doi: 10.1016/j.culher.2017.01.013.

[34] Bali, A. and Singh, S. A Review on the Strategies and Techniques of Image Segmentation. *2015 Fifth International Conference on Advanced Computing & Communication Technologies*, pages 113–120, 2015.

[35] Ghosh, C., Majumder, S., Ray, S., Datta, S., Nath, S., and Ghosh, C. Different EDGE Detection Techniques: A Review.

[36] Azzopardi, G. and Petkov, N. Contour Detection by CORF Operator. In *Artificial Neural Networks and Machine Learning – ICANN 2012*, pages 395–402. Springer Berlin Heidelberg, 2012. ISBN 978-3-642-33269-2.

[37] Ilea, D. and Whelan, P. Image segmentation based on the integration of colour-texture descriptors - A review. *Pattern Recognit.*, 44:2479–2501, 2011.

[38] Rohani, N., Salvant, J., Bahaadini, S., Cossairt, O., Walton, M., and Katsaggelos, A. Automatic pigment identification on roman Egyptian paintings by using sparse

modeling of hyperspectral images. pages 2111–2115, 08 2016. doi: 10.1109/EUSIPCO.2016.7760621.

[39] Tarabalka, Y., Chanussot, J., and Benediktsson, J. Segmentation and classification of hyperspectral images using watershed transformation. *Pattern Recognition*, 43: 2367–2379, 07 2010. doi: 10.1016/j.patcog.2010.01.016.

[40] Wu, H., Wu, Y., Zhang, S., Li, P., and Wen, Z. Cartoon image segmentation based on improved SLIC superpixels and adaptive region propagation merging. *2016 IEEE International Conference on Signal and Image Processing (ICSIP)*, pages 277–281, 2016.

[41] Wang, X., Ma, P., and Zhao, J. Brain tumor CT image segmentation based on SLIC0 superpixels. *2016 9th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, pages 427–431, 2016.

[42] Dong, R., Wang, B., Li, S., Zhou, Z., Li, S., and Wang, Z. Interactive image segmentation with color and texture information by region merging. *2016 Chinese Control and Decision Conference (CCDC)*, pages 777–783, 2016.

[43] Li, L., Yao, J., Tu, J., Lu, X., Li, K., and Liu, Y. Edge-based split-and-merge superpixel segmentation. In *2015 IEEE International Conference on Information and Automation*, pages 970–975, 2015. doi: 10.1109/ICInfA.2015.7279427.

[44] Hsu, C.-Y. and Ding, J.-J. Efficient image segmentation algorithm using SLIC superpixels and boundary-focused region merging. In *2013 9th International Conference on Information, Communications Signal Processing*, pages 1–5, 2013. doi: 10.1109/ICICS.2013.6782861.

[45] Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., and Susstrunk, S. SLIC Superpixels Compared to State-of-the-Art Superpixel Methods. *IEEE transactions on pattern analysis and machine intelligence*, 34, 05 2012. doi: 10.1109/TPAMI.2012.120.

[46] Balas, C., Epitropou, G., Tsapras, A., and Hadjinicolaou, N. Hyperspectral imaging and spectral classification for pigment identification and mapping in paintings by El Greco and his workshop. *Multimedia Tools and Applications*, 77, 04 2018. doi: 10.1007/s11042-017-5564-2.

[47] Liu, X. and Yang, C. A Kernel Spectral Angle Mapper algorithm for remote sensing image classification. *Proceedings of the 2013 6th International Congress on Image*

and Signal Processing, CISP 2013, 2:814–818, 12 2013. doi: 10.1109/CISP.2013.6745277.

[48] Chen, M.-Y., Huang, Y.-B., Chang, S.-P., and Ouhyoung, M. Prediction Model for Semitransparent Watercolor Pigment Mixtures Using Deep Learning with a Dataset of Transmittance and Reflectance, 2019.

[49] de Carvalho Júnior, O. and Meneses, P. Spectral Correlation Mapper (SCM): An Improvement on the Spectral Angle Mapper (SAM). Online, January 2000.

[50] Bioucas-Dias, J. A variable splitting augmented Lagrangian approach to linear spectral unmixing. volume 1, pages 1 – 4, 09 2009. doi: 10.1109/WHISPERS.2009.5289072.

[51] Dessein, A., Cont, A., and Lemaitre, G. Real-time Polyphonic Music Transcription with Non-negative Matrix Factorization and Beta-divergence. *Proceedings of the 11th International Society for Music Information Retrieval Conference, ISMIR 2010*, pages 489–494, 08 2010.

[52] Winsor & Newton, Composition and Permanence. Website. URL https://bit.ly/3yCaTxi. Accessed February 2021.

[53] *Specim FX10 / FX10e - User Guide 1.5*. Specim Spectral Imaging.

[54] Bartyzel, K. Adaptive Kuwahara filter. *Signal, Image and Video Processing*, 10:1–8, 07 2015. doi: 10.1007/s11760-015-0791-3.

[55] Yuhas, R. H., Goetz, A., and Boardman, J. Discrimination among semi-arid landscape endmembers using the Spectral Angle Mapper (SAM) algorithm. 1992.

[56] Ding, J., Li, X., and Huang, L. A novel method for spectral similarity measure by fusing shape and amplitude features. *Journal of Engineering Science and Technology Review*, 8:172–179, 12 2015. doi: 10.25103/jestr.085.22.

[57] Geng, X., Ji, L., and Sun, V. Nonnegative matrix factorization based unmixing for principal component transformed hyperspectral data. *Frontiers of Information Technology & Electronic Engineering*, 17, 04 2016. doi: 10.1631/FITEE.1600028.

[58] Guo, B. Hyperspectral Image Classification via Matching Absorption Features. *IEEE Access*, PP:1–1, 09 2019. doi: 10.1109/ACCESS.2019.2940268.

[59] Yokoya, N., Chan, J., and Segl, K. Potential of Resolution-Enhanced Hyperspectral Data for Mineral Mapping Using Simulated EnMAP and Sentinel-2 Images. *Remote. Sens.*, 8:172, 2016.

[60] Rahiche, A. and Cheriet, M. Forgery Detection in Hyperspectral Document Images using Graph Orthogonal Nonnegative Matrix Factorization. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2823–2831, 2020.

[61] Arthur, D. and Vassilvitskii, S. K-Means++: The Advantages of Careful Seeding. page 1027–1035. Society for Industrial and Applied Mathematics, 2007.

[62] Noris, G., Hornung, A., Sumner, R. W., Simmons, M., and Gross, M. Topology-Driven Vectorization of Clean Line Drawings. *ACM Trans. Graph.*, 32(1), 2013. doi: 10.1145/2421636.2421640.

[63] Ricciardi, P., Delaney, J., Glinsman, L., Thoury, M., Facini, M., and de la Rie, R. Use of visible and infrared reflectance and luminescence imaging spectroscopy to study illuminated manuscripts: Pigment identification and visualization of underdrawings. *Proceedings of SPIE - The International Society for Optical Engineering*, 7391, 07 2009. doi: 10.1117/12.827415.

[64] Magro, N., Bonnici, A., and Cristina, S. Hyperspectral image segmentation for paint analysis. In *IEEE International Conference on Image Processing*, 2021. Accepted for Publication.

Appendix A

Histograms for Pixel Offset

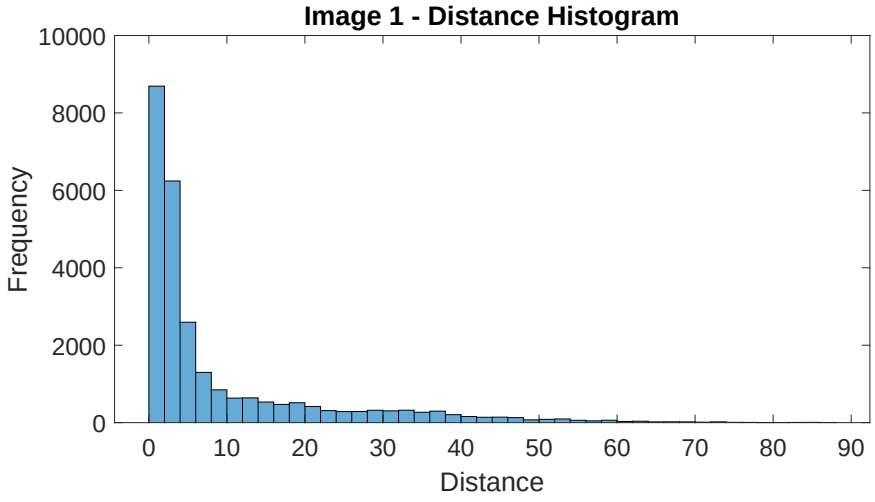

(a)

FIGURE A.1: Pixel Distance Probability Density plot for Image 1.

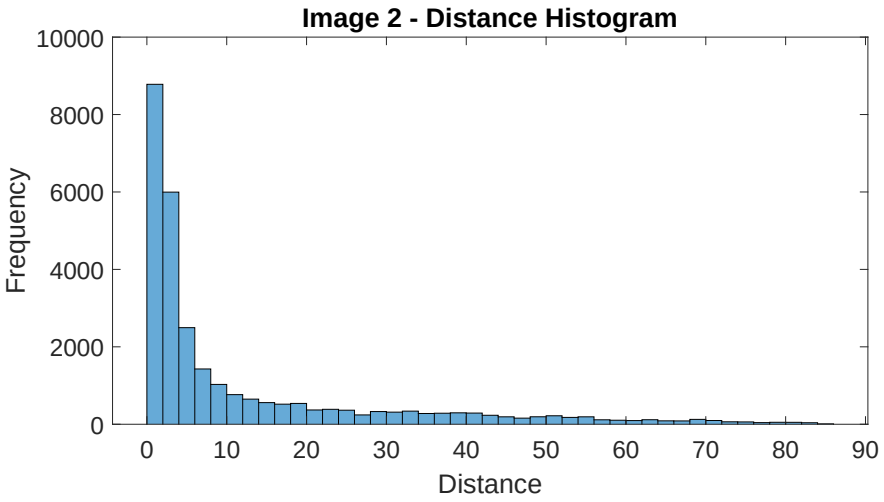

(a)

FIGURE A.2: Pixel Distance Probability Density plot for Image 2.

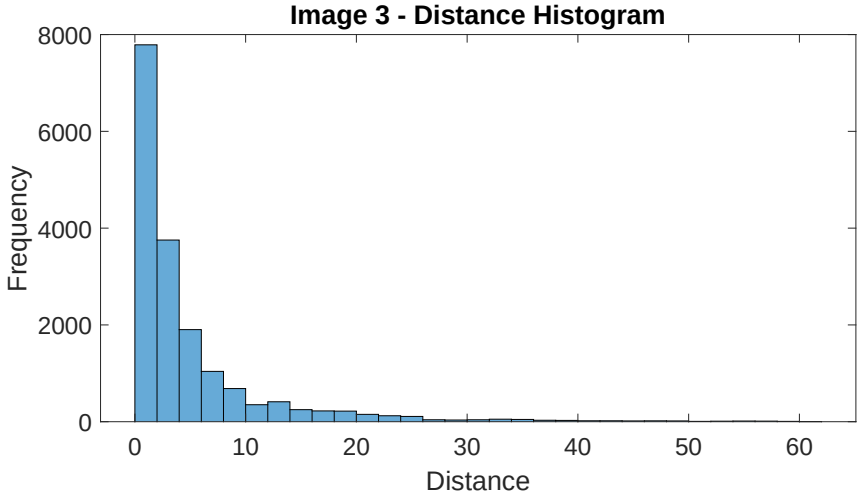

(a)

FIGURE A.3: Pixel Distance Probability Density plot for Image 3.

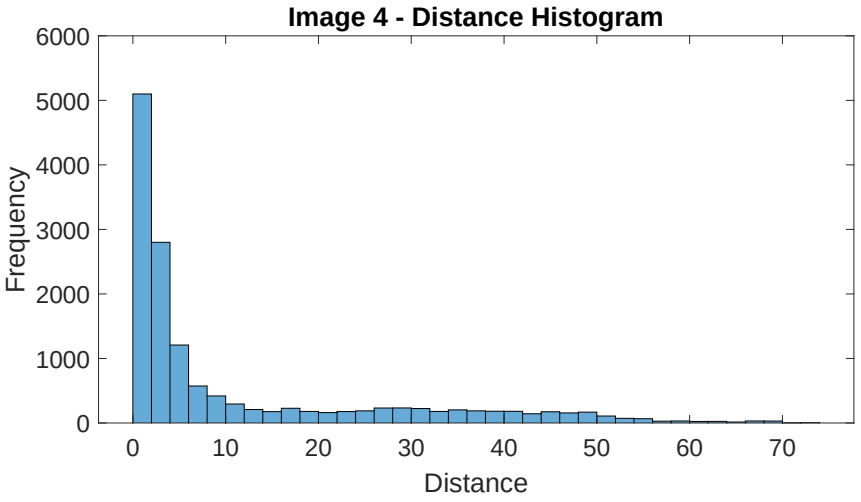

(a)

FIGURE A.4: Pixel Distance Probability Density plot for Image 4.

Appendix B

Classification results for other images

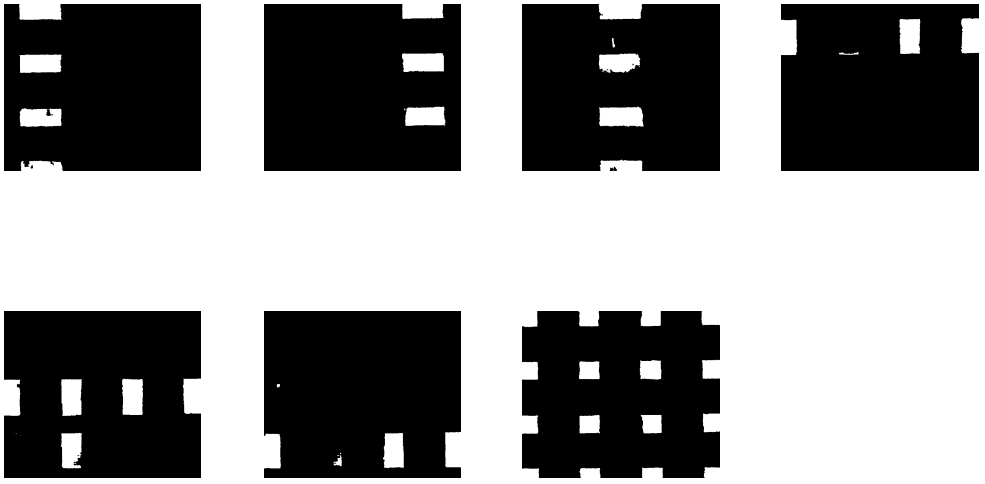

FIGURE B.1: Hierarchical Paint Analysis result for Image 2, for stage 1 of the algorithm which regards to separating individual regions from mixed regions.

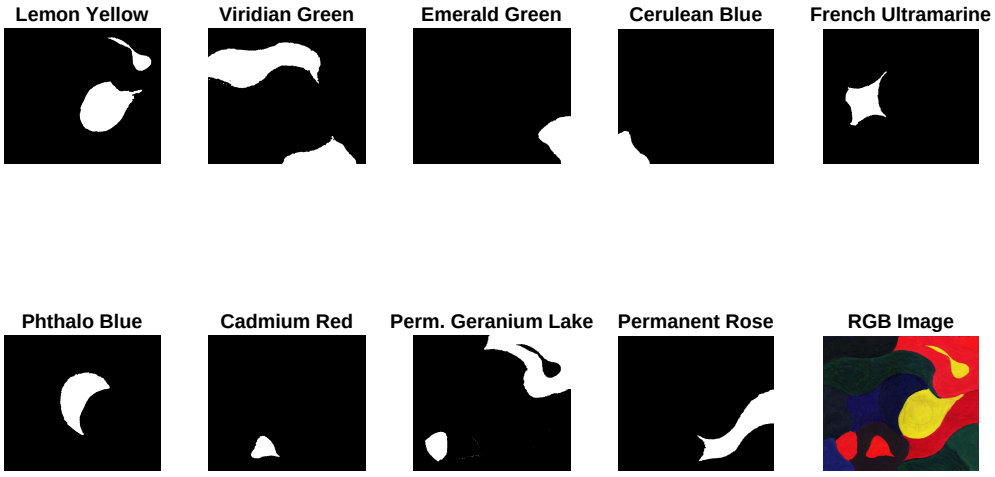

FIGURE B.2: Hierarchical Paint Analysis result for Image 3, for stage 1 of the algorithm which regards to separating individual regions from mixed regions.

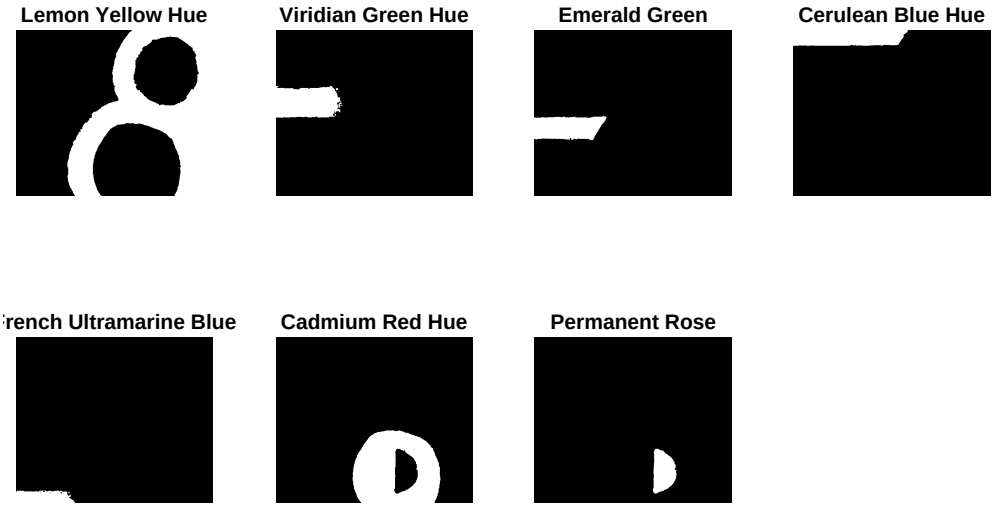

FIGURE B.3: Hierarchical Paint Analysis result for Image 4, for stage 1 of the algorithm which regards to separating individual regions from mixed regions.

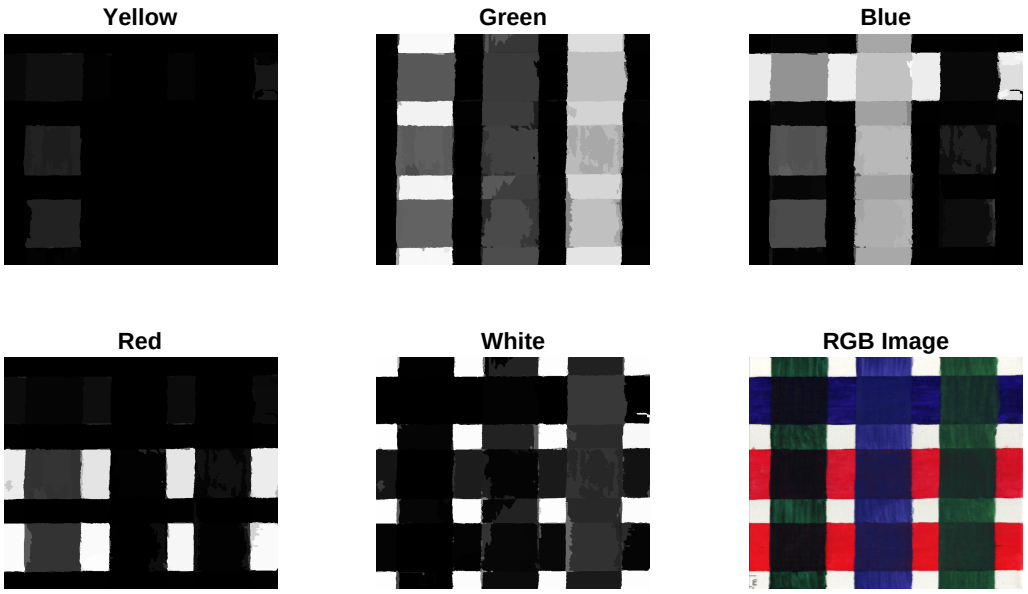

FIGURE B.4: NNMF Classification result of Image 1 when separating the paint regions by paint class.

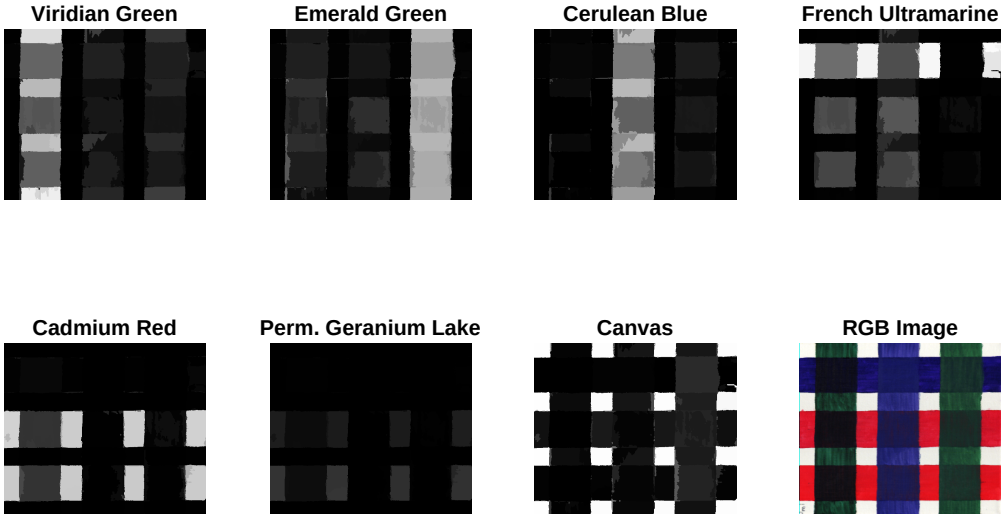

FIGURE B.5: NNMF Classification result of Image 1 when separating the paint regions by individual paints.

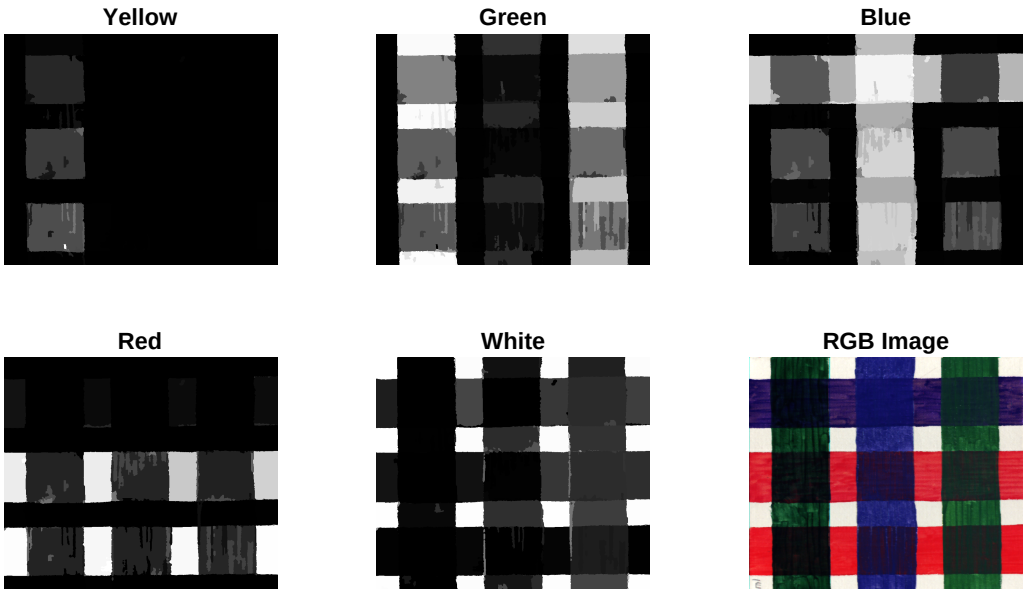

FIGURE B.6: NNMF Classification result of Image 2 when separating the paint regions by paint class.

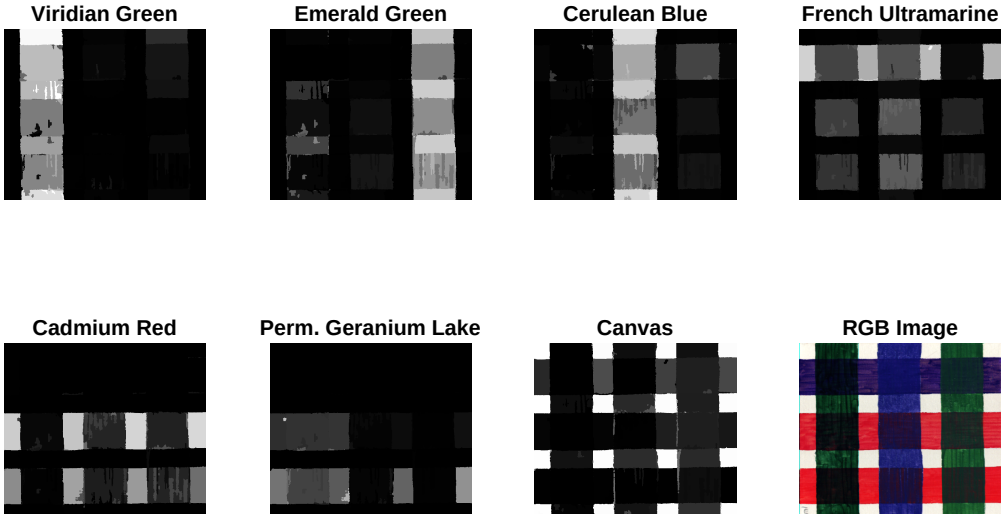

FIGURE B.7: NNMF Classification result of Image 2 when separating the paint regions by individual paints.

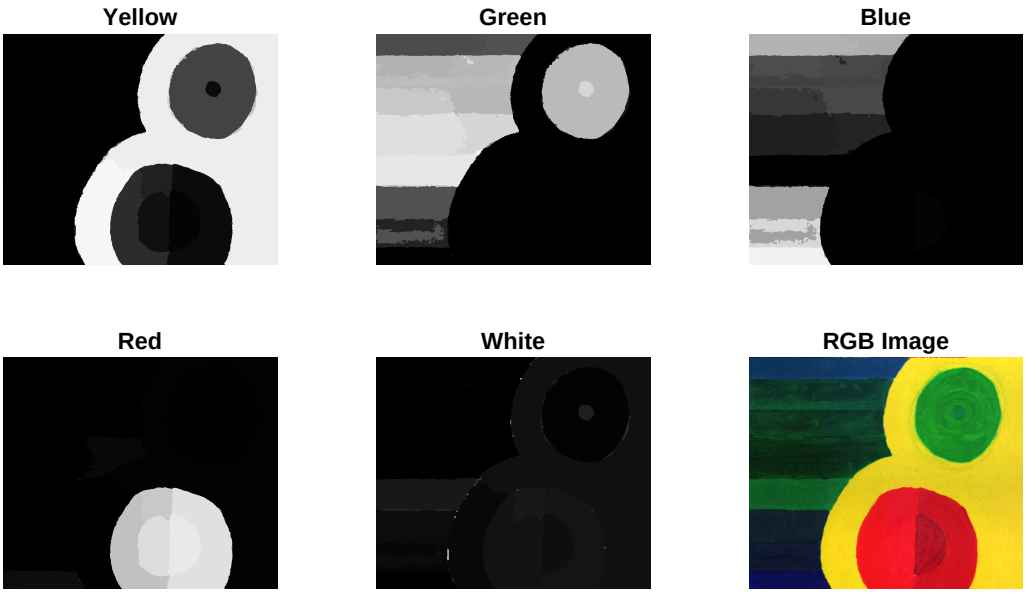

FIGURE B.8: NNMF Classification result of Image 4 when separating the paint regions by paint class.

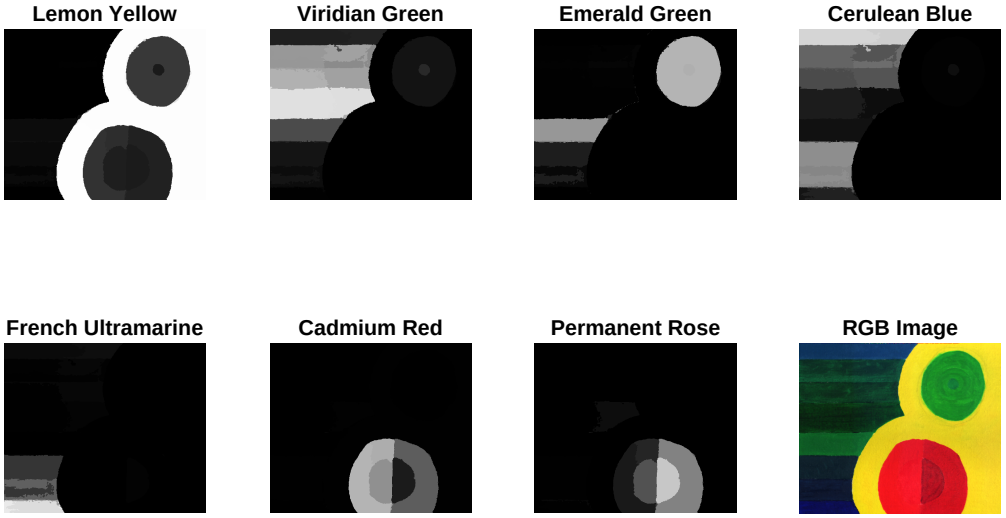

FIGURE B.9: NNMF Classification result of Image 4 when separating the paint regions by individual paints.

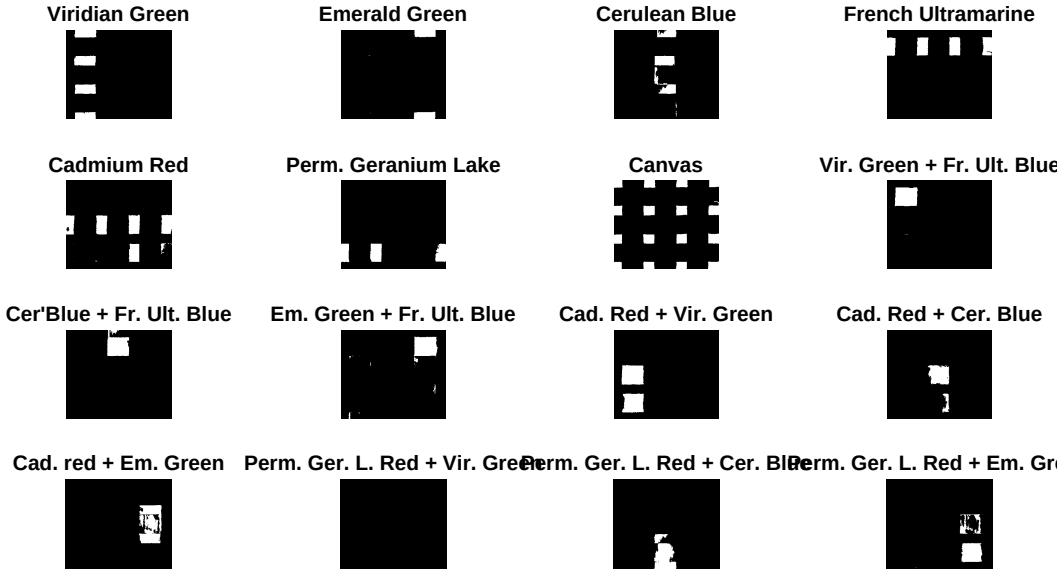

FIGURE B.10: Classification result of Image 1 for KSCM. A mixed paint reference is provided in the reference library. class.

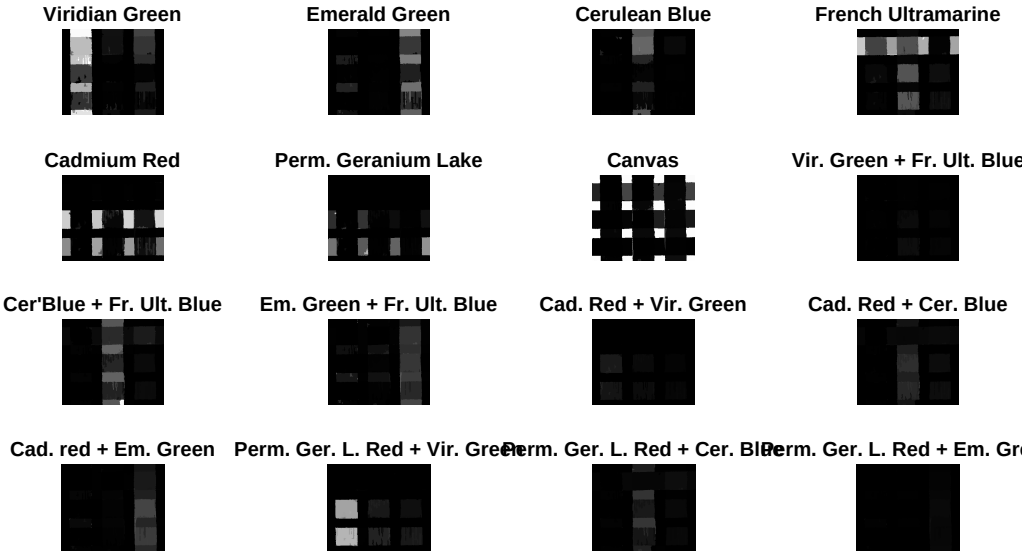

FIGURE B.11: Classification result of Image 1 for NNMF. A mixed paint reference is provided in the reference library.

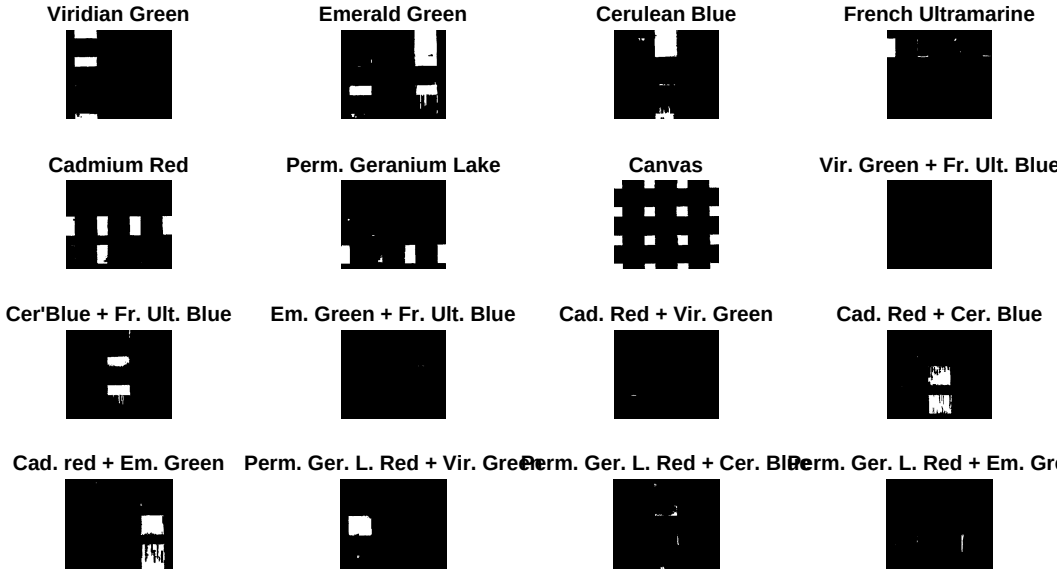

FIGURE B.12: Classification result of Image 2 for KSCM. A mixed paint reference is provided in the reference library. class.

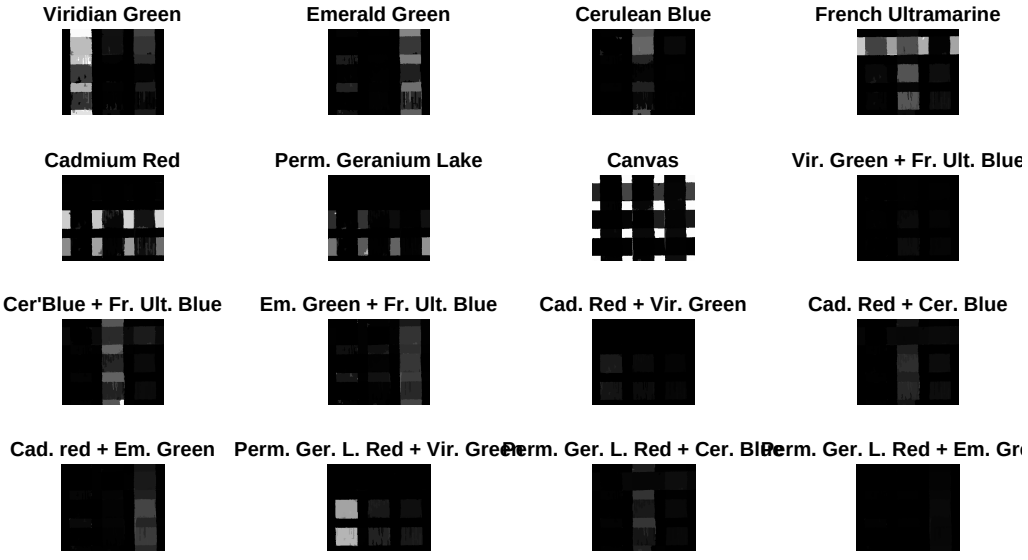

FIGURE B.13: Classification result of Image 2 for NNMF. A mixed paint reference is provided in the reference library.

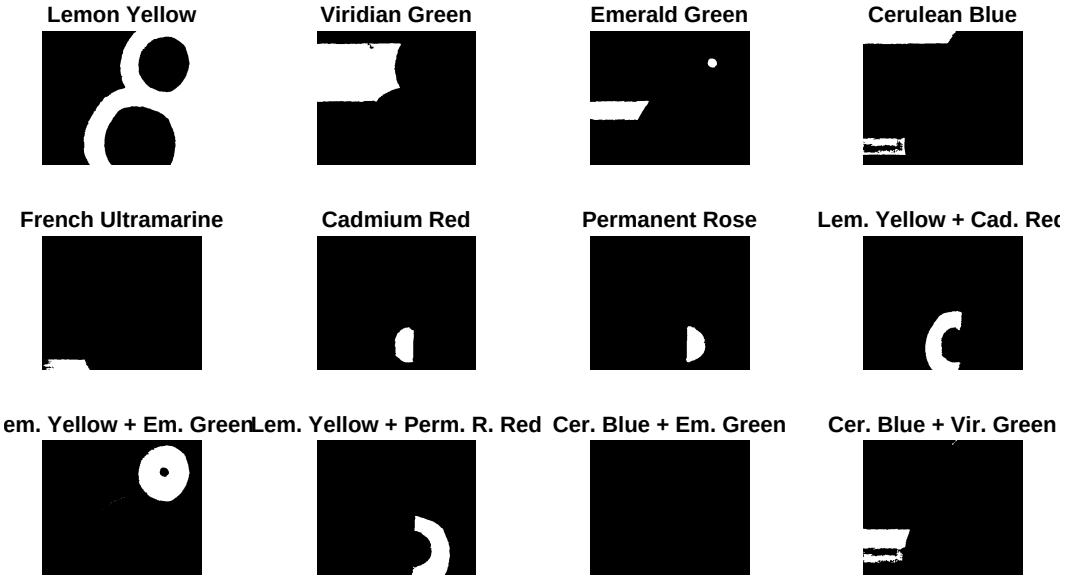

FIGURE B.14: Classification result of Image 4 for KSCM. A mixed paint reference is provided in the reference library. class.

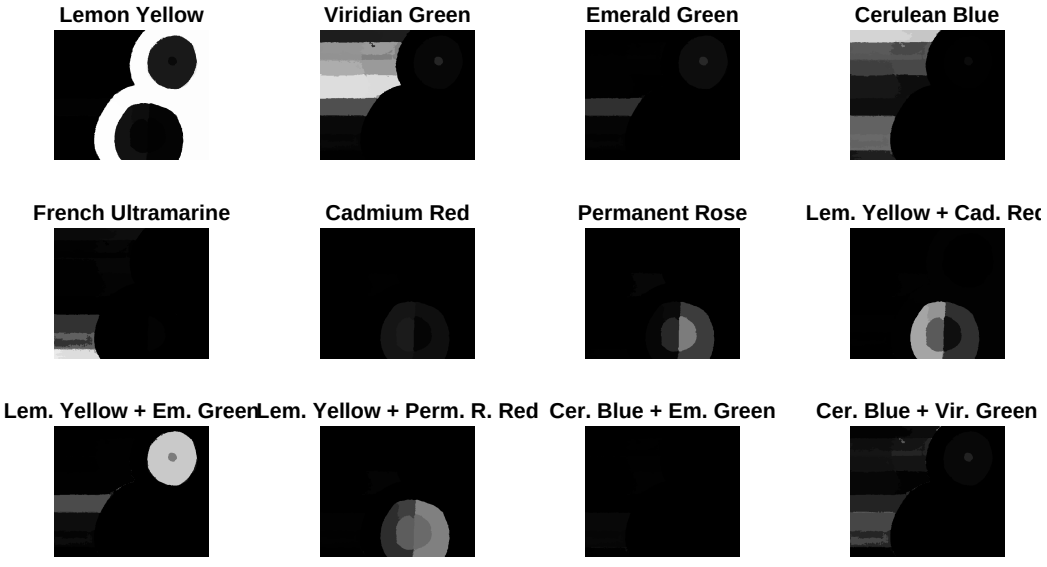

FIGURE B.15: Classification result of Image 4 for NNMF. A mixed paint reference is provided in the reference library.