



J. Belder

Investigating palettes used in the painting of Rembrandt's *Man with the red baret* using X-ray fluorescence scanning and t-distributed stochastic neighbour embedding



Investigating palettes used in the painting of
Rembrandt's *Man with the red baret* using
X-ray fluorescence scanning and
t-distributed stochastic neighbour embedding

By

J. Belder

in partial fulfilment of the requirements for the degree of

Master of Science

In Materials Science and Engineering

at the Delft University of Technology,
to be defended publicly on 20 April 2018 at 3:00 PM.

Student Number:	4001672	
Supervisors:	Prof. dr. J. Dik Prof. dr. B.J. Thijsse	
Thesis committee:	Prof. dr. J. Dik Prof. dr. B.J. Thijsse Prof. dr. R.G. Erdmann	TU Delft TU Delft Rijksmuseum

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Preface

For the last year I've been working on this thesis by having a look at the *Man with the red baret*. I'm grateful for the opportunity to combine art, programming and material sciences into a single research topic. Finishing this research also marks the end of my time as student at the TU Delft. For the last eight and a half years I've been working to reach this point. Sometimes I was focussed on studying, sometimes I had my attention elsewhere. Looking back, I realise that there is much I learned, on a technical level, professional level as well as personal level. I am proud of the work that is accomplished here and I look forward to moving into the next chapter of my life.

I couldn't have made this thesis without the help and support of others. I'd like to thank Joris Dik for his enthusiasm in the proceeding of the research and all the suggestions for things to discover. Also a huge thanks to Barend Thijsse for all the encouragement to do sound statistics and his sharp eye for details that otherwise would've been missed. Finally I'd like to thank Rob Erdmann for taking place in the graduation committee.

Also I'm grateful for all the support and love from my parents. Both during the work on my master thesis as well as throughout my whole life. I wouldn't have gotten here if it were not for you. A huge thank you!

Furthermore I want to thank the folks at Insyde for having the flexibility to combine work and studying. I also want to thank all my friends, roomies, fellow-students, co-workers for all the *gezelligheid* and good times that kept me outside the bubble of pure studying.

Last and certainly the most I wish to thank God, whom by, amazing grace, chose to be our heavenly Father. Thank You, Father, for Your amazing grace and unfailing love. Thank You for creating and sustaining everything there is. Thank You for enabling us discover, uncover, explore and learn new things about the world around us. I can only wonder what You'll reveal to us next.

*Joan Belder
Delft, 6 April 2018*

Table of contents

Preface	IV
Abstract.....	1
Introduction	2
Contents	2
Background.....	4
Painting process	4
X-ray fluorescence.....	5
Structure of painting	9
Methods	12
Used data in this work	12
Fitting	12
t-Distributed stochastic neighbour embedding	14
Chi-squared test over sections	16
Results and discussion.....	18
The sum spectrum.....	18
Fitting.....	18
Sectioning using t-SNE.....	22
Sections within the painting	28
Conclusions	32
Recommendations.....	33
References	34
Appendix A: Extended table of fitting parameters	36
Appendix B: Elemental distribution maps.....	37
Appendix C: extended χ^2 -test result values.....	44
Appendix D: extended report per section	45

Abstract

It is assumed that 17th century paintings were painted by a very systematic approach which involved the use of very specific palettes. For every pictorial unit, such as a skin, a very specific palette was created. In this case the palette would then only be used to paint the skin and nothing else. The result is that for each pictorial unit a separate palette is used. This assumption can be verified by characterizing the all the paint mixtures present in a painting. Paint derives its colour from the pigments present within them. Many pigments used in the 17th century contain characteristic heavy elements. The usage macro X-ray fluorescence (XRF) scanning is gaining popularity as a non-destructive method for determining elements present in a painting. By the creation of elemental maps the pigments can be identified and localized in a painting. These maps already allow for visualizing many interesting features such as retouches, inpaintings and even hidden layers behind a painting. However the usage of the XRF data have never been used for identifying and visualizing the palettes used while painting.

In this work it is attempted to identify and visualize the palettes used in the *man with the red baret* by Rembrandt van Rijn, by an in-depth analysis of raw XRF data. The raw data was model fitted using a simulated annealing algorithm to retrieve spectral line intensities in each point of the painting. The fitted spectral intensities were then visualized using t-distributed stochastic neighbour embedding (t-SNE). With help of the t-SNE visualisation the painting is divided into 24 sections. The suggestion is that these 24 sections each contain a different and unique paint mixture. This suggestion is further investigated by the use of the two-sample χ^2 -test between the sections' spectral composition. The combination of model fitting alongside with t-SNE on a XRF spectra can indeed visualize different paint mixtures and give strong suggestion of the different palettes used in *man with the red baret*.

Introduction

The *Man with the red baret* (Dutch: Man met de rode muts) is a painting attributed to Rembrandt van Rijn (1606 – 1669) or his studio and was painted around 1660 [1]. The painting is shown in figure 1. It has been suggested that in this time period a painter worked with a structured palette for a specific type of painting [2]. This means that for example when a portrait was being painted that a specific portrait palette was prepared using a standard recipe prior to painting. A recipe would consist of the mixing of a number of pigments, such as lead-white and vermilion with linseed oil, in order to achieve certain effects in the painting. This systematic approach of painting has been suggested, but never been investigated by analytical and statistical means. In this thesis the *Man with the red baret* by Rembrandt will be investigated for the use of standard palettes, by determining the paint mixtures present in the painting. By determining the palettes used, we hope to gain insight in the process of the creation of the artwork itself. Also it enables us to find retouches and inpaintings that are not distinguishable by investigating individual elemental distributions maps.

The investigation of palettes is done by an in-depth analysis of the raw data from a X-ray fluorescence (XRF) scan [3]. The analysis is done in several steps. First model fitting is performed, to gain pure elemental data so that contaminations such as background radiation are filtered out. Secondly, t-distributed stochastic neighbour embedding was on the elemental distribution maps to visualize elemental groups in the painting [4]. This makes it possible to recognize various *sections* in the painting in which the elemental composition is different. In our case 24 of these will be discovered. These *sections* of the painting allow to take a closer look at paint compositions and give insight in the construction process of the artwork. The uniqueness of the sections is then verified by using the χ^2 -test.

This work only investigates the use for a palette of a single painting. However the methods used are applicable on other paintings as well. In fact the methods used are not solely restricted to the use of XRF-scanning data, but also visual images, hyperspectral data and the results from other imaging techniques or any combination of these can be used to visualize palettes used in the artwork. Even within the limited scope of only XRF scanning there is much to be discovered for this single painting. We can only wonder what the methods used will uncover for other works of art.

Contents

The following subjects will be discussed in this thesis:

The *background* of the research will be covered. The painting practice of the 17th century will be covered. This will also explain the significance of palettes for a work of art. The second background topic is X-Ray fluorescence. We will cover some of the XRF basic topics, and more importantly how XRF interacts with a painting itself. To have a deeper understanding of these interactions, the structure will also briefly be discussed.

Then the *methods* used in this work are discussed. In this section all the details of the raw XRF data, model fitting, t-SNE and further statistical analysis of these painting sections will be covered. Also an example will be given on the working of t-SNE, to further clarify what this technique does.

Thirdly the *results* be discussed. The results for the individual steps of model fitting, t-SNE, and χ^2 -test will be discussed. Then we will get into what all these results tell us about the artwork itself.

Finally it's time to wrap things up by summarizing the most important *conclusions* and give some *recommendations* about further research.



Figure 1: Rembrandt van Rijn, Man with red baret, ca. 1660, Museum Boijmans Van Beuningen, Rotterdam, The Netherlands.

Background

Painting process

In figure 2 a print is shown of the workplace of an Renaissance era painting workshop [2]. Although this is a highly idealized image, it gives us an idea how a painting is constructed. It can be seen that while the master painter is making the large painting, his assistants (at the top right) are grinding and dissolving the pigments into oil. Also there is an assistant preparing the palette for the master painter in the front of the image. This print gives an impression of painters worked in the 17th century. Most importantly we see that the paint was specifically prepared for what the painter was working on. Also that using all the oils a specific paint mixture was prepared on the palette before applying to the painting itself. The palette is then limited to paint a single a single figurative unit (e.g. the skin of a man). Hence the painter would paint this figurative unit and when something else had to be painted the painter would then have to switch his palette. This means that each figurative unit is associated with a specific palette but also that all units would be painted one at a time. Besides this print this approach to painting is also suggested by written sources. The Dutch painter Willem Beurs describes in his book *De Groote Waereld in het Kleen Geschildert* (The big world painted in small) how to paint a white horse [5]:

One paints the illuminated side using white, light ochre and black, with pure white for the highlights; light ochre is recommended for the intermediate colour, and it is advisable there to be rather sparing with the white. For the shadow, black and light ochre must be mixed together with a little white; the reflection under the belly should be mostly light ochre, with sparing use of black and white. The hooves can sometimes be painted with black, white and light ochre, with a touch of vermillion; and sometimes with black, white and umber. The colour of the nose is the same as that of the hoofs. But as for the eyes: the pupil should be painted with bone black, and the rest with umber, black and white.

If one were to paint a white horse using this recipe, only a limited number of pigments would be needed: white lead, yellow ochre, vermillion, umber and bone black. If a painting depicts a white horse (and possible many other things), then it can be expected that where the horse is painted only the 5 mentioned pigments will be found. Thus a horse is defined by these 5 pigments, while another figurative unit might be defined by another set of pigments. In the end the painting is then composed of a set of figurative units which all have their distinctive pigments.

Van der Wetering suggest that this workflow used by the 17th century painters are due to the technical constraints the artists had to work with [2]. The paint itself could not be stored and for each paint a specific recipe had to be used upon creation [6]. Only up until the invention of tube paint around 1840 the technical advances allowed to paint whole painting simultaneously [2] [7].

If indeed paintings were made using a single palette for each pictorial entity, this should have results on the final painting: each figurative will contain only a limited number of pigments; a figurative unit will contain a unique mixture of pigments; some pigments are also clearly absent; and pigments should be present locally and not smeared out over the whole painting. It is already known that even the composition of the binding medium (linseed oil, walnut oil) changes over the various pictorial entities of paintings [6]. Finally it is likely that the composition of pigments is not exact and varies from one palette to the other, even if visually two colours might look the same. However the presence of palettes has, as of today, never been studied on a large number of points in a single painting using technical analytical means.



Figure 2: Jan Baptiste Collaert, Color Olivi, ca. 1590. After Johannes Stradanus

X-ray fluorescence

X-ray fluorescence is a non-destructive technique to determine the elemental contents of a sample. This is done by irradiating the target with hard (high energetic) X-rays, while simultaneously measuring the response X-rays that come from the target. The result is the elemental composition of the sample for the elements $Z > 12$ can be determined. When a thin X-ray beam scans over the surface of an 2d object the elemental composition can be determined at every point. This makes it possible to create elemental an maps of surface of the object. For scanning an object the X-ray source is typically a synchrotron beam or an X-ray tube. A Synchrotron has a very strong incident beam and allows for short dwell times and thus (relatively) fast scanning. However since a synchrotron requires are large facility, the current portable solution for paintings uses an X-ray tube.

Atoms can have various responses to X-rays, in XRF the measured response are the lower energy X-rays that are the result of ionization of the atoms in the target. An incident X-ray photon is absorbed by one of the inner electrons in an atom which is then ejected from its orbit out of the atom. The newly created vacancy is then filled by one of the electrons from an outer shell thereby emitting an characteristic X-ray photon. A schematic of two such cases is given in figure 3 [8].




Figure 3: Schematic of the process of creating a characteristic x-ray fluorescence radiation. In this case the $K\alpha$ and the $K\beta$ radiation of an atom are shown. [8]

The shells around an atom are commonly referred to by a letter, starting with K for the innermost shell, then L for the secondary shell, M for the third shell etc. Subshells are usually numbered with L_1 being closest to the nucleus while L_3 is the furthest away from the nucleus. In order to successfully create an photon transition from the L₃-shell the K-shell the incident beam must have enough energy to successfully knock out the K-shell electron out of its orbit into the vacuum. Thus an energy of $E_{in} > E_k$ is required. The energy of this characteristic X-ray photon equals to the difference between the shell energies:

$$E_{L_3 \rightarrow K} = E_k - E_L \quad (1)$$

Siegbahn notation

For energy transitions from one shell to the other the Siegbahn notation is commonly used in literature. This work will also use the Siegbahn notation. The most common transition $L_3 \rightarrow K$ is for example referred to as the $K\alpha_1$ radiation of an element. An overview of the most common transitions in Siegbahn notation is given in figure 4 [9], which is a more precise version of figure 3. An alternative to the Siegbahn notation is the IUPAC notation which explicitly denotes the transition between to electron shells [10]. The XRF setup used for analysing the painting in this study does not have the energy resolution required to practically measure the $K\alpha_1$ and the $K\alpha_2$, the $L\alpha_1$ and $L\alpha_2$ radiation separately. These transitions (especially in the case with lighter elements) tend to have similar energies and therefore these peaks are colloquially referred to with the $K\alpha$ transition or the $L\alpha$ without the number.




Figure 4: Overview of the most common electron shell transitions in XRF and the associated Siegbahn notation. Taken from the X-Ray data booklet [9].

Chemical bonds

For light elements (especially sulphur in the case of paintings) the valence electrons might be involved in the characteristic X-ray generation process. The valence state of an atom affects the outer electron energies. When the energy is different also a slight shift of the expected characteristic energy should be observed. In some special cases the valence state of an atom can then even be determined using XRF [11]. However not only the outer electrons but also the inner electrons are also affected by the valence state of the atom, counteracting the shift of the characteristic energy [12]. For sulphur the difference in energy is < 10 eV which too low to be resolved by the detector system used in this work.

Considerations for scanning

When doing XRF measurements there is always some background radiation detected. This is due to the Compton scattering effect. An X-ray photon can collide inelastically within the material losing some of its energy. These X-rays that have lost energy also arrive at the detector and provide a baseline of background radiation even if no characteristic radiation is expected in the energy range.



Figure 5: Compton scattering intensity map between 16.1 and 16.7 keV of 'Man with red baret'. It can be seen that the stretcher and the supporting beams give rise to a higher background radiation level. Even the easel which was used during the scanning of the painting can be observed. Secondly it can be seen that some sections are shielded due to lead-white presence in the paint layer.

Also the shielding of radiation should be considered. Heavy elements absorb more X-ray radiation than light elements. The intensity of background radiation also varies over the painting. In figure 5 a portion of the background radiation of the painting is shown. It can be seen that the the lead-white intense areas (figure 6) shield off the background radiation. Also the supporting frame appears lighter, because a photon passes more molecules increasing the chance for inelastic scattering.

The third consideration is that high energetic X-rays have a deeper penetration than low energetic X-rays. This means that for example lead-white paint at the surface will absorb more X-rays from deeper layers than a section that has no lead-white paint at the surface. The recorded intensity of an element behind a layer of lead-white than a layer that is behind a lighter element. Even if concentration and layer sizes are kept the same. In a similar fashion low energy radiation can easily be detected when originating from a surface layer, but might be difficult to observe when it's originating from a deeper found layer in the painting. This fact could in theory be exploited to gain insight in the depth of an element. Comparing the relative intensity of the characteristic radiation of multiple spectral lines from a single element gives a measure of the shielding for that element. For example when comparing the high energetic Pb Ly elemental map to a low energetic Pb M α elemental map in figure 6 some differences can be observed. Pb Ly can originate from deeper layers than Pb M α can, therefore it could be concluded that parts in which Pb Ly but not Pb M α is an indication that the lead is to be found in a deeper layer. Thus it's important to keep in mind that not all elements are easily detected even when only present in deeper layer. Also the intensity is not a purely dependent on the concentration of a given element and should not be used as such.



Figure 6: Comparison of the high energetic Pb Ly and low energetic Pb Ma elemental map. In the Pb Ly map the hat clearly shows the presence of lead, while in the Pb Ma no lead is detected. This suggests that the lead is to be found in deeper layers.

Structure of painting

The man with the red baret is an easel painting and this typically this type of painting consists of multiple layers which can be roughly divided into the following four: the support, the ground layer, then several paint layers (for simplicity underpaintings are also counted among these) and finally a layer of varnish. Figure 7 shows a rough schematic of the various layers of a painting. These layers consist of various materials which in turn do affect how X-rays interact with these materials. The various layers will be further discussed in this chapter.




Figure 7: A schematic cross section of a painting. Generally four layers do exist: the support, the ground layer, the paint layers and on top the varnish.

Support material

Usually the painting is supported by a piece of stretched linen canvas on a frame or a wooden panel. In the case of *Man with the red baret* the support consists of a piece of canvas on a frame. Additionally the frame is supported with some wooden bars, which is common for larger paintings. The support is usually treated with *size*. This typically consists of diluted animal glue and prevents the binder material (oil in an oil painting) to be absorbed by the support. Absorption of binder can be detrimental for both the painting layers on top of the support as well for the support material itself and is therefore unwanted [13].

All the components that make up the support are usually organic and consists of light elements that are practically undetectable by XRF. The variation in thickness of the background layer increases the chance for Compton-scattering and thus increases the background radiation. Also layers on top of the support can shield off a certain level of Compton scattering from the detector. Hence the support structure of the painting can be well visualized using XRF as seen in figure 5.

Ground layer

In order to provide a smooth and clear background for painting a ground layer is applied to the canvas. For historic paintings this layer typically consists of gesso: a mixture of chalk, animal glue and sometimes also a diluted pigment, in order to provide a baseline for the rest to be drawn upon. A way to temper the colour of the baseline sometimes a transparent paint layer is applied to establish a baseline colour this is called an imprimatura. An underdrawing made from charcoal or a dark paint can be present on the painting which can be both on top of the imprimatura layer as well as on the bottom of the imprimatura layer. This provides a rough sketch in what the painter envisions the final result to be of the artwork.













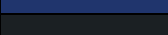
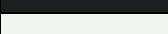

Paint layers

On top of the base layers the actual paint is applied. The painting usually consists of an underpainting establishing a base outline of the final painting and after that various paint layers are applied to achieve the final result. The paint consists of a binder medium and a pigment to give its colour. In oil paintings the binder medium consists of oils such as linseed oil or walnut oil.

Around 20 pigments were commonly used for 17th century paintings, and most of these pigments contain traces of elements whose mass is larger than sulphur ($Z = 18$), and are therefore detectable by XRF. Solely organic pigments are currently undetectable by the means of XRF. This means that using only XRF techniques not every type of paint can be accounted for.

In table 1 a non-exhaustive list is given for XRF analysis for the *Man with the red baret*. The list is compiled based on the elements that have been found using the X-ray data from this study. Rembrandt is known to have used more pigments in his works [14]. This means that pigments that are composed of elements which were not found are omitted. For example no Cu was found on the painting, hence Malachite ($\text{Cu}_2\text{CO}_3(\text{OH})_2$) was left out.

Table 1: A non-exhaustive list of pigments that should be considered when looking at the 'Man with the red baret'. The colour is an indication of the colour of the pigment, and is not by any means an attempt to provide an exact representation.

Pigment	Characteristic XRF elements	Colour	Remarks
Raw sienna	Fe, Mn		
Burnt sienna	Fe, Mn		
Raw umber	Fe, Mn		
Burnt umber	Fe, Mn		
Red ochre	Fe		
Yellow ochre	Fe		
Lead-white	Pb		Contains many impurities [15].
Lead-tin-yellow	Pb, Sn		
Vermillion	Hg, S		
Orpiment	As, S		
Realgar	As, S		
Smalt	Ni, Co, K		Tends to degrade on old paintings.
Bone-black	Ca		
Titanium-white	Ti		Used after 1900
Zinc-white	Zn		Used after 1850

Varnish

The final layer of the painting is the varnish. The varnish layer serves two purposes. First and foremost varnish serves as a protective barrier between the air and the paint layer. This helps degradation due to pollutants in the air, moisture and dirt. Secondly the varnish ensures that a glossy look on a painting will be uniform, since various pigments may vary in how matte or glossy they appear.

The varnish layer is usually based on a resin such as mastic or dammar. The varnish layer tends to gradually become yellow over time. Especially mastic, which was commonly used in the 17th century, is prone to this yellowing process [16] [17]. Figure 8 shows the effect of yellowing varnish over a long period of time. As a standard maintenance and restoration procedure the varnish layer of a painting is removed and a new varnish layer is then applied to the painting [16]. It is unlikely that any restored work from the 17th century still has its original varnish.

Both classical and modern varnishes usually consist of an organic resin and are therefore not detectable by an XRF scanner. The varnish layer is therefore not considered when doing the data processing in this work.



Figure 8: Image by Philip Mould. Varnish tends to become more yellow due to ageing effects. This image shows the effect of 200 years of yellowing varnish compared to areas where the varnish layer is removed.

Methods

Used data in this work

The data was measured by a Jetstream m6 XRF scanner [18]. The painting is scanned in two parts: a top part and a bottom part. The areas that are covered by each scan is shown in figure 9. The top part of the painting is scanned with a grid of 1292 × 941 pixels, where each pixel contains an energy spectrum. The bottom part was scanned in a grid of 1233 × 940 pixels. The dwell time during scanning on each pixel was 100 ms. A spectrum at each pixel consist of the photon counts over 4096 channels in the energy range of -0.9557 keV to 39.9943 keV.

In order to improve the signal to noise ratio at the cost of spatial resolution the data was 3×3 binned. This was done by first cropping the file to 1290 × 939 pixels discarding the bottom right. This step is necessary in order to have both the horizontal and vertical dimensions divisible by 3. Subsequently a binning was done by summing the channel counts over a 3×3 area. A group of 3×3 pixels in the spectrum is thus reduced to a single '9pixel'. After binning the top part has the dimension has 430 × 313 9pixels. It's important to note that at every 9pixel there still is a full energy spectrum of recorded photon counts.




Figure 9: Visual representation of the scanned area. The pink outline is the top part. The green the bottom part.

Fitting

Once the data from the scanner is retrieved and condensed to 9pixels, one would like to retrieve the elements present in each location of the painting. This can be done by comparing the measured spectrum with theoretically expected spectra for specific elemental emission lines. To identify the elements present in the painting a sum spectrum was made. This was done by summing the counts of each specific channel over all the pixels. Peaks can be observed in the sum spectrum that correspond to characteristic XRF-emissions. These elemental peaks are then subsequently model fitted. A list of fitted peaks is given in table 2. Some elemental peaks were purposely omitted. For example Ar is present in the air, but not in the painting so gives no insight in the painting. Additionally Zn and Ti were left out because these elements are associated with restorations and thus give no insight in the use of palettes by the original painter. Additionally some peaks would provide little extra value since the same

element has been fitted already with another spectral line. An example would be Pb L₁ which is to be expected to be similar to Pb L_α, Pb L_β and Pb L_γ.

After the sum spectrum was fitted, each individual *9pixel* also needs to be fitted for every spectral line. It is unfeasible to fine-tune the fit for every *9pixel* out of the roughly 130,000 pixels present. Therefore it is important that for each fit there is a proper initial guess to succeed in fitting. The fitting results of the sum spectrum can be used to create an initial guess for the fitting process for an individual pixel. After fitting the sum spectrum it is known what the peak height is for all pixels. Also the number of pixels is known. This makes it possible to calculate the average peak height per pixel. The same applies for the intensity of the background radiation. The initial guess parameters for each individual *9pixel* fit is exactly this estimation of the average of all parameters when all pixels are fitted.

Table 2: Spectral lines that were fitted to the sum spectrum. For the energy values the energy of the α₁ and β₁ peak is taken from literature [19].

Spectral line	Energy (keV)
As Kβ	11.727
Ca Kα	3.692
Co Kα	6.930
Co Kβ	7.649
Cu Kα	8.048
Fe Kα	6.404
Fe Kβ	7.058
Hg Lα	9.999
Hg Lβ	11.823
K Kα	3.314
Mn Kα	5.899
Ni Kα	7.478
Pb Lα & As Kα ¹	10.555 (Pb) & 10.544 (As)
Pb Lβ	12.618
Pb Lγ	14.769
Pb Mα	2.345
Pb Mβ	2.442
S Kα	2.308
Sn Kα	25.272

The fitting model for both the sum spectrum and each individual spectrum was in the following form.

$$C(x) = a + bx + H_1 e^{-\frac{1}{2} \left(\frac{x-\mu_1}{w_1} \right)^2} + H_2 e^{-\frac{1}{2} \left(\frac{x-\mu_2}{w_2} \right)^2} + \dots \quad (2)$$

In which x is the energy in keV of a channel (data point) in the fit, $C(x)$ the estimated number of counts and $a + bx$ is the a linear background portion over the energy range. Then for each peak n within the given energy range a gaussian curve is added, where μ_n is the energy of the spectral line in keV and w_n the standard deviation of the n^{th} Gaussian curve. Independent fits have shown that the measured peak shapes are near-perfectly Gaussian. Both a and b were unconstrained during fitting, while the constraints $H_n > 0$ and $0.03 < w_n < 0.15$ keV were added applied. A full overview of all the parameters used for fitting can be found in appendix A. The fitting process itself is to find the optimal solution by minimizing χ^2 with the simulated annealing method [20]. This is given by the formula:

$$\chi^2 = \sum_x d(x)^2 \quad (3)$$

In which $d(x)$ is the fit residue for a given channel. This residue is defined as:

¹ The energy of Pb-Lα As-Kα are in practice inseparable by the model fitting algorithm used due to the proximity of both peaks. Thus this energy was fitted as if it were a single spectral line.

$$d(x) = \frac{C(x) - M(x)}{\sigma(x)} \quad (4)$$

In which $C(x)$ is the estimated energy and $M(x)$ is the actual measured counts per pixel at energy x , and $\sigma(x)$ is the standard deviation of counting in the xrf scanner. Due to the fact that photon counting obeys Poisson statistics and needs a small correction for background noise this can be estimated as:

$$\sigma(x) = \begin{cases} \sqrt{M(x)}, & M(x) \geq 10 \\ \sqrt{M(x)} + 1 - 0.1M(x), & M(x) < 10 \end{cases} \quad (5)$$

The quality of the fit was additionally verified by calculating the Inverse Renormalized Cyclic Durbin-Watson-Hendrikse statistic (ψ^2) [21] using the formula:

$$\psi^2 = \frac{\sum_x d(x)^2}{0.5 \sum_x d(x + c_w) - d(x)^2} \quad (6)$$

In this formula the c_w is the energy width of a single channel. This is equal to 0.01 keV. A good fit has the following properties:

$$\begin{aligned} E[\chi^2] &= \nu, & \text{std}(\chi^2) &= \frac{1}{\sqrt{2\nu}} \\ E[\psi^2] &= 1, & \text{std}(\psi^2) &= \frac{1}{\sqrt{\nu + 2}} \end{aligned} \quad (7)$$

In which ν is the number of degrees of freedom in the fit, $E[\chi^2]$ the expected value of χ^2 , $E[\psi^2]$ the expected value of ψ^2 , $\text{std}(\chi^2)$ the standard deviation over χ^2 and $\text{std}(\psi^2)$ the standard deviation over ψ^2 . The degree of freedom can be calculated from the amount of channels in the fit N_x and the amount of variables to be optimized during fitting k .

$$\nu = N_x - k \quad (8)$$

When fitted the area of the under the Gaussian peak can be calculated by:

$$I_n = \sqrt{2\pi} H_n \sigma_n \approx 2.5 H_n \sigma_n \quad (9)$$

The area under the curve is an indication of the number photons. This area should be used when relating different intensities with each other.

t-Distributed stochastic neighbour embedding

A single pixel on the painting can also be seen as a point in a high dimensional space with as axes the N_u spectral lines and the location of that point is the intensity of each fitted fluorescence line along the respective axis. In other words each pixel has measured intensity As K β , Ca K α , Co K α etc. In a system with an axis for As K β and an axis for Ca K α etc. the pixel is can be found in a coordinate in a 19 dimensional space. It is impossible to visualize 19 dimensions in a scatter plot. Therefore a form of dimension reduction or embedding is required. For this dimension reduction t-distributed stochastic neighbour embedding (t-SNE) was used [4]. t-SNE attempts to reduce the 19 dimensions to a 2d plane by placing pixels in the 2d plane in such a way that pixels that two pixels that have a small Euclidean distances in the high dimensional space also have a small distance on the 2d plane. In this work this 2d plane that is constructed the t-SNE image. This embedding is performed with the *scikit-learn* machine learning library. Specifically the Barnes-Hut implementation of t-SNE [22] [23] was used. One of the required parameters for t-SNE is the perplexity. In t-SNE the perplexity can be seen as an estimation for the number of similar points for each point [4] [24]. The t-SNE was performed at various perplexities ranging from 20 – 60. The perplexity of 35 is found to be the easiest to interpret hence this perplexity is chosen for further analysis.

The other t-SNE parameters that are used are a learning rate of 50 and a maximum of 5000 iterations. The rest of the parameters are set to the *scikit-learn* default settings.

t-SNE example

In figure 10 an example is shown when performing the t-SNE over the rgb-channels of an image of a painting by Piet Mondriaan. Since there are three separate channels (red, green and blue) in this case $N_u = 3$. In this example it can be seen that separate clusters are formed after transformation. Also each cluster is composed of pixels that have similar colours. In the Rembrandt case instead of the rgb-value of each pixel, the intensity for each spectral line is considered. The idea is that the result is to be similar. In the ideal scenario there will also be clusters formed. Like each cluster of pixels in the painting of Mondriaan has a similar colour, each cluster of points in the Rembrandt painting would then have a similar spectral energy composition.




Figure 10: When performing t-SNE on the rgb-channels of the image on the left (Mondriaan, Piet. "Composition with Red, Yellow, Blue and Black." 1921) it can be seen that after transformation multiple clusters have appeared, each having a specific colour. This colour of a point in the t-SNE image corresponds with a pixel in the original image. This simple example shows that t-SNE is extremely successful in clustering points with a similar rgb-value. For the Rembrandt painting in this work the fluorescence energies instead of the rgb values are used.

Selection and sections

The main use of t-SNE is to assist in dividing the painting into different groups of pixels. Each group of pixels will be called a *section* throughout the rest of the work. It should be mentioned that a single pixel can only be part of single *section*. Each *section* or group of pixels has been assigned a number in order to compare them with each other. The choice of sections is strongly related to the clusters of points visible within the t-SNE image. The process of defining the sections based on the observed clusters in the t-SNE image is explained in detail in the chapter *sectioning using t-SNE*.

Pre-processing

Prior to t-SNE analysis all N_u distributions for spectral line energy u over all pixels are normalized so that each distribution has the following properties:

$$\mu_{norm,u} = 0; S_{norm,u} = 1 \quad (10)$$

In which $\mu_{norm,l}$ is the normalized mean of all the intensities of a single fluorescence line, whereas $S_{norm,l}$ is the standard deviation of all the intensities of a single fluorescence line. In order to achieve this, each spectral intensity per pixel is normalized by point the following formula:

$$I_{norm,p} = \frac{I_p - \mu_u}{S_u} \quad (11)$$

In which $I_{norm,p}$ is the normalized intensity at energy u for a single pixel p , I_p is the calculated intensity due to fitting this pixel, μ_u is the average of I_p for all pixels and S_u is the standard deviation of the I_p values over all pixels.

t-SNE limitations

t-SNE is a very powerful visualization technique for finding interesting features in high dimensional data. However it can be difficult to find the most useful result due to the large number of parameters involved. Changing some parameters can yield quite different images. Especially the perplexity has a large impact in the image. The amount of clustering can change drastically depending if extreme perplexities are chosen. However the results generally seemed quite stable with ranging perplexities from 25 to 50.

t-SNE attempts to get to the optimal image by the means of gradient descent. This makes it very likely that t-SNE image will show a local optimum of the arrangement of points. This means that sometimes two similar clusters are not merged in the final image, while in a global optimum they would have. Also if the starting points for the 2d representation are chosen randomly, running t-SNE multiple times will result in different visualisations. It is even advisable to run it multiple times. This way the best visualisation can then be selected as the one to be used for further analysis.

Third the t-SNE image should be interpreted very carefully. For example if the distance between two clusters in t-SNE is unimportant. If two clusters are separate in the visualisation this means that in the high dimensional space there is also a separate cluster. However if two clusters are far apart in the visualisation it does not have to mean that these clusters are far apart in high dimensional space. Also from the shape of a cluster in the visualisation not much can be derived. There is however one exception: if a cluster in the visualisation is connected by string of points to another cluster this means that it is likely that also in the high dimensional data two clusters are connected by string of points between these clusters.

However with all limitations in mind t-SNE is very effective in finding clusters in high dimensional space. One of the most powerful features of t-SNE is that the shape of the cluster in the original high dimensional space does not matter to successfully recognized as cluster in the low dimensional space. Hence t-SNE can be very successful in finding different spectral compositions with varying intensity. This means that it can finding different elemental compositions which in turn allows for identifying different paint mixtures in the painting. It allows for finding specific mixtures that one might not easily see by looking at each individual spectral map alone. Hence it's a great method to both identify and localize the paint mixtures and used. Once each paint mixture is identified and localized, they can be assessed individually: is this mixture specific for an inpainting or is it part of the original painting layer? Is this mixture the result of stacking of paint layers or is it a mixture by itself? What does it mean that a mixture is different from another mixture but have the same colour visually?

Chi-squared test over sections

There is also the question to answer whether two sections are also really different, or whether t-SNE is giving false positives. For this the Pearson's chi-square test can be used. This is because the average spectral composition of an *section* (a group of pixels) can also be seen as categorical data. By using the chi-squared test, two sections can be compared against each other for significant difference. The two sample test is the test to use in this case [20]. For each section the mean per spectral line $\mu_{i,u}$ and the sample standard deviation per spectral line $S_{i,u}$ is calculated. Then both were corrected for varying intensities by dividing through the sum of the means of a section:

$$\mu_i^{sum} = \sum_u \mu_{i,u} \quad (12)$$

$$\mu'_{i,u} = \frac{\mu_{i,u}}{\mu_i^{sum}}, \quad S'_{i,u} = \frac{S_{i,u}}{\mu_i^{sum}} \quad (13)$$

To calculate the $\hat{\chi}^2$ -statistic between section i and section j the following formula was used:

$$\hat{\chi}_{i,j}^2 = \sum_u \left(\frac{\mu'_{i,u} - \mu'_{j,u}}{\sqrt{S'^2_{i,u} + S'^2_{j,u}}} \right)^2 = \sum_u \frac{(\mu'_{i,u} - \mu'_{j,u})^2}{S'^2_{i,u} + S'^2_{j,u}} \quad (14)$$

The number of degrees of freedom ν for the chi-square probability density function $Q(\chi^2|\nu)$ distribution is given by:

$$\nu = N_u - 1 = 19 - 1 = 18 \quad (15)$$

In which N_u is the number of spectral lines considered for this test. For this work 19 spectral lines were considered in all occasions: therefore it can be stated that ν is equal to 18 and therefore a p -value can be calculated that section i and section j have a different mean elemental composition. The hypothesis that section i and j are different in this respect is accepted if the calculated $\hat{\chi}^2$ is in the 5% tail in $Q(\chi^2|\nu)$, or put in formula form:

$$p = P(\chi^2 > \hat{\chi}^2) < 0.05 \quad (16)$$

Since the number of degrees of freedom ν is known. The critical $\hat{\chi}^2$ can be calculated:

$$\hat{\chi}_{crit}^2 = \text{cdf}(\chi^2) > 0.95 \approx 30.14 \quad (17)$$

in which the cdf function is the cumulative density function of the χ_{18}^2 -distribution. Therefore for verification if two of the sections can be simplified to the following condition: $\hat{\chi}^2 > 30.14$.

Two remarks must be made: first the 5% uncertainty level is often in agreement with common statistical practice. Secondly the specific form chosen for the uncertainty of the mean line intensity in a given section $S_{i,u}$ in equation 14 is not the only one conceivable. The one used here expresses our choice that we want to assess the compositional difference between two sections *relative* to the compositional variation *within* each of the two sections.

The chi-square statistic was calculated between every possible combination of sections in the painting. From the statistically different sections are the conclusions drawn in this work.

Results and discussion

The sum spectrum

Figure 11 shows the sum spectrum and it can be seen that indeed many different spectral lines are distinguishable. The most important peaks are all fitted. Some peaks are deliberately ignored such as the Ar K α peaks due to the fact that this peak does not contain relevant information for the painting itself. Also some peaks in the higher energy ranges are instrumental artefacts in the sense that they have double the energy of a lower energy peak. For example the peak at 21 keV might be attributed to Cd K α , however there are not any other spectral lines visible for Cd, and this peak is located at 2×10.5 (Pb L α) keV and show also a strong relation to other Pb peaks.

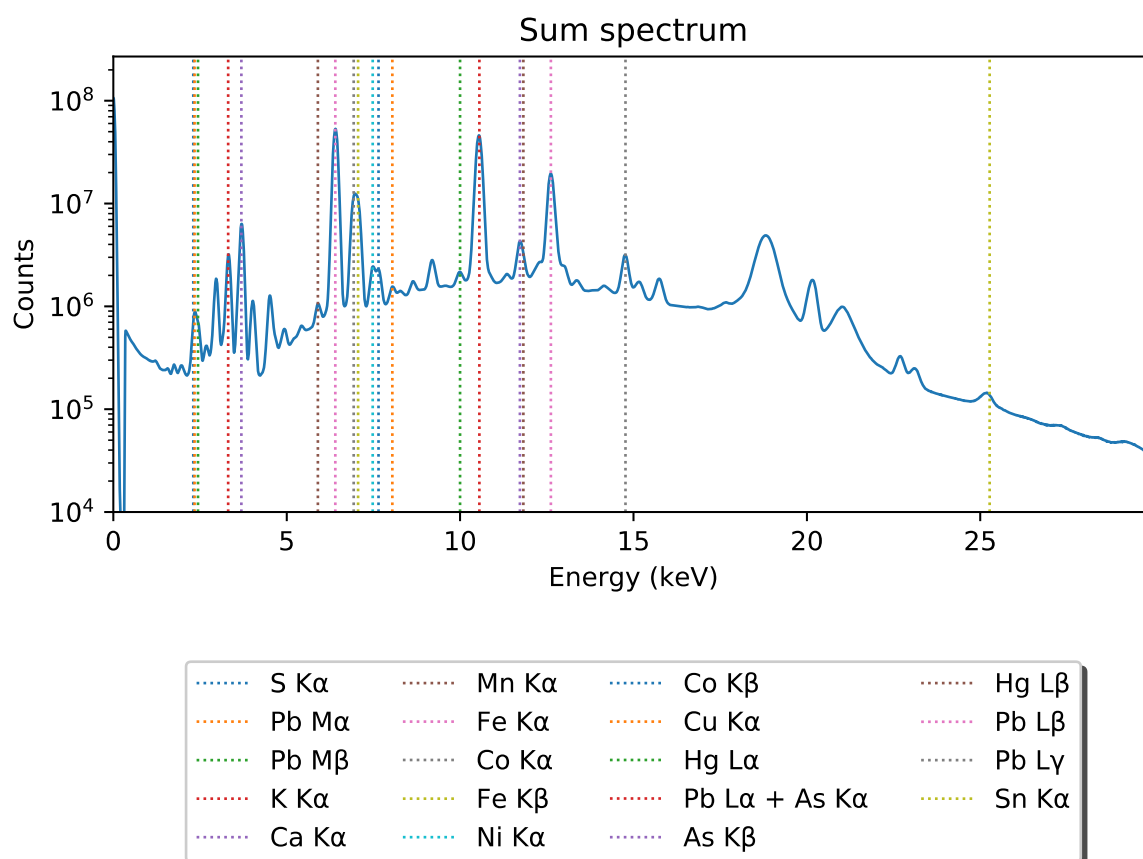


Figure 11: Sum spectrum of the top part of the painting. All fitted lines are marked in the chart. For peak energies, see Table 2.

Fitting

By model fitting the data not only the background can be distinguished from the signal, also overlapping signals can be separated and further statistical analysis can be performed. These capabilities are to be considered separately. First of all the consistency of the fitting can be considered and it can be determined if the fitted data is consistent with itself. This is a method to verify if the fitting was performed correctly. Secondly fitting overlapping peaks allows for detecting elements that are not distinguishable otherwise. This is the case for sulphur which will be discussed later on. Finally, statistical analysis will make it possible to validate several aspects of the measurement and analysis procedures.

Consistency

The main reason for consistency checks is the following: the fitting algorithm applied here for non-linear least-squares fitting (simulated annealing) does not guarantee that the minimum χ^2 will be found. It executes a series of iterative steps using random numbers, and even repeated under identical conditions may yield non-identical results. This is why none of the fits is guaranteed to be the absolute best possible, although the method will end up very close to this best fit in the majority of the cases.

To gain confidence in that the fitting is performed correctly is to check the internal consistency of the fitting result. A very simple, yet powerful, method is to examine if two different spectral lines of the same element, therefore of the same atoms, correlate with each other. In our case it makes sense to check if the Pb L β intensities from the individual pixels correlate with the Pb L γ intensities. In figure 12 a comparison is shown between Pb L β and Pb L γ , and also between Pb M α , Pb M β , Fe K α and Fe K β and Co K α and Co K β . It can be seen that these expected correlations do indeed exist. It should be noted that the correlation is not fully linear in at least one case (Pb L). It remains unclear what the exact cause is for this behaviour, though it should be noted that these lines are at different energies and X-ray penetration mechanics, such as shielding and level of absorption, should not be ruled out. The noise in the correlation plots is due to counting statistics and fitting uncertainty.




Figure 12: Comparison of the pixel intensities of the Pb L, Pb M, Fe K and Co K lines. The units is counts \cdot keV. It can be seen that there is a strong correlation between these pairs of line intensities. It should also be noted that the correlation is not fully linear in the Pb L case.

Other internal correlations can also be considered to check for consistency: for example it is known that smalt contains both cobalt and nickel therefore it is expected that these two elements are correlated with each other if substantial quantities of smalt are present in the painting. In figure 13 this is clearly the case, which also confirms that fitting is consistent among different elements.




Figure 13: Scatter plot of Co K α and Ni K α pixel intensities. Both elements are associated with the pigment smalt.

Finally the resulting χ^2 and the ψ^2 values are an indication whether the fitting was successful. As previously mentioned in equation 7. In table 3 an overview of the average found χ^2 and ψ^2 is given. It can be seen that for most fitted energy ranges the average is indeed in the range of 0.7 – 1.3, and thus relatively close to the expected 1. There are however two exceptions.

The first exception is the $\sqrt{\chi^2/\nu}$ -statistic of Sn K α peak, which is below 0.7. The low χ^2 -statistic means for Sn K α is most likely due to over-estimation of the uncertainty as given in equation 5. Since the Sn K α peak has the lowest intensity of all fitted peaks (see figure 11), this is likely an effect of the correction factor in equation 5.

The second exception is average ψ^2 -value of Pb L α + As K α peak which is above 1.3. This suggests the model is not fully correct. In this case indeed both Pb L α and As K α were attempted to be fit into a single peak. However if both Pb L α and As K α peaks were fitted separately the chi-squared and psi-squared statistics would even increase further, which is an indication that either that model is not fully correct or that the fitting algorithm (simulated annealing) is unable to find the best solution. Finally this also shows the merit of also using the ψ^2 -statistic. It enables for finding errors within the fitting model, that might otherwise be hidden when using the χ^2 -statistic.

Table 3: The average ($\sqrt{\chi^2/\nu}$) and the average ψ^2 of each pixel after fitting. Both statistics have an expected value of 1.

Fit peaks	Energy Range (keV)		$\sqrt{\chi^2/\nu}$	ψ^2
S K α , Pb M α , Pb M β	2.124	2.564	0.864	0.961
K K α , Ar K β	3.114	3.474	0.989	0.925
Ca K α_1 , Ca K α_2	3.504	3.824	0.982	0.913
Mn K α	5.704	6.034	0.936	0.945
Fe K α , Mn K β	6.044	6.574	1.071	1.043
Co K α , Fe K β	6.674	7.284	1.056	1.011
Ni K α , Co K β	7.294	7.844	1.017	0.951
Cu K α	7.854	8.174	0.999	0.939
Hg L α_1 , Hg L α_2	9.704	10.154	1.064	0.977
Pb L α + As K α	10.164	10.974	1.200	1.393
Hg L β As K β	11.494	12.004	1.063	0.959
Pb L β	12.334	12.874	1.100	1.118
Pb L γ	14.504	14.974	1.028	0.980
Sn K α	24.004	25.984	0.510	1.091

Determining sulphur

One of the main advantages of fitting is that it allows for properly splitting overlapping peaks in the spectrum. One particular peak of interest is the S K α line which is the only detectable sulphur line (2.308 keV). However the energy range between 2.2 and 2.6 keV is dominated by the Pb M α and Pb M β lines, as can be seen in figure 14. Although also the S K β lines could also be present in this spectrum, the relative intensity of S K α may be so low compared to and Pb M α and Pb M β lines that usually this line cannot be observed [9].

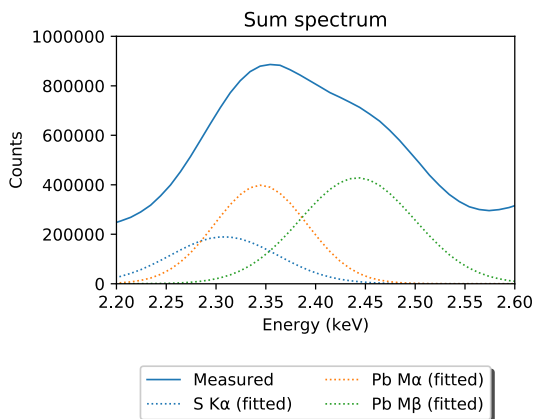


Figure 14: Part of the sum spectrum showing the S K α , Pb M α , Pb M β fitted peak contributions.. The individual peaks are not visible, but a combination of these peaks is visible with predominantly Pb visible. Note that for clarity the background Which can also be seen on the corresponding distribution map of the energies between 2.2 and 2.5 keV on the right side. The lead-white heavy parts such as the face and the collar of the man predominantly show in this distribution.

Without fitting techniques the sulphur would be undetectable. However when fitting sulphur is detectable and can be distinguished from the Pb rich parts. This can also be seen in Figure 15, where the sulphur distribution is clearly different from the Pb distribution.



Figure 15: Comparison of the fitted $S\ K\alpha$ and the $Pb\ M\alpha$ peaks of the top part of the paintings. It can be seen that the S heavy parts are corresponding to the reddish parts of the painting itself which would correspond to the pigments vermillion (HgS) and Realgar (As_4S_4), while the Pb rich parts are predominantly caused by lead-white at the surface.

Sectioning using t-SNE

Once the fitting is done, there is need to consider all the fitted elemental maps simultaneously. Using t-SNE all the 19 elemental maps can be visualized in a special way in a special way explained earlier. Figure 16 shows the result performed on the top part of the painting. The transformation shows the individual pixels moved in the original image in such a way that pixels with a comparable spectral composition are close to one another. No recolouring of the pixels was applied. Spectral lines correspond with elements, which in turn correspond with pigments. So it can be expected that a similar spectral composition also means a similar pigment composition. If similar pigments are grouped within the t-SNE image it can be expected that also some of the clusters in this image have their own distinct colour. In figure 16 this is the case. For example most of the lighter face colours can be seen in the top of the t-SNE image.




Figure 16: t-SNE transformation result based on the 19 fitted elemental intensities of all pixels. The points in the t-SNE plot are coloured based on the colour from the photograph in which the point was originally located on the painting.

Clustering

In the t-SNE image certain clusters and subclusters of points can be recognized, similar to the clusters of the example with Mondriaan in figure 10. It should be noted that the position of clusters does *not* contain information about the elemental composition. Even when two clusters are next to one another they might have a totally different composition, while two clusters that might be more similar can be at the opposite sides of the t-SNE image. In t-SNE only the distances between all points are considered. When an image is rotated the distance between points does not change. This means that when a t-SNE image can also be rotated freely.

In fact the exact x,y-coordinate in which the point is located in the t-SNE plot does not even matter. What only matters to what clusters can be seen in the t-SNE plot and to a lesser extent the shape of the clusters that are formed.

So one can derive very little from the t-SNE image only by itself. But it does provide an excellent starting point for deriving palettes. A cluster of points in the t-SNE image is expected to have a similar spectral composition. A cluster thus provides a starting point for what might have been a palette. Remember that a similar spectral composition also implies a similar elemental composition, which in turn implies a similar mixture of pigments in the paint. A cluster of points corresponds to a number of pixels in the original painting. And it should be remembered that each pixel contains the spectral intensity of 19 fitted peaks as well as a position in the original painting. So when a cluster of points is considered in the t-SNE image, one is also considering a selection of pixels within the original painting. From a selection of pixels an average and standard deviation of every spectral peak intensity can be calculated. Also of a selection the pixel location in the original image can be highlighted.

For each cluster that is recognized in the t-SNE image the pixel intensity average and standard deviations for elemental each peak is calculated as well as the original locations are mapped. It is found that some clusters are quite similar both in spectral intensity composition, and also represent a coherent pictorial entity. For example two clusters with similar spectral intensity composition are both depicting the hair of the man in the painting. Also multiple clusters are found that would compose the background of the scene in the painting. When this is the case the clusters are assumed to be of the same palette.

A cluster that has been found unique or a set of clusters that are assumed to be the same will be called *sections* throughout this report. This means that a *section* can consist of *multiple* clusters in the t-SNE image. Secondly a cluster that has been recognized in the t-SNE is can never be a part of multiple sections. However in most cases a *section* corresponds to exactly a single cluster within the t-SNE image.

Using this method a total of 24 sections are recognized in the painting. For convenience these sections have been numbered from 1 to 24. In figure 17 an overview is given of all sections. Additionally in appendix D an extended report is given per section. In the numbering process of sections the choice is made to give the sections that consist of many pixels a low number, while the sections that have a low pixel count have been assigned a high number.

In figure 17 it can be seen that for the sections 1 – 11 that have been mapped back to their original position the major features of the painting can be recognized. This in itself is a promising result: it means that t-SNE can be used as a visualisation technique for the high dimensional data in a sense that it does recognize important features. The original position or colour of a pixel is not considered by t-SNE; these features are recognized purely by the spectral peak intensities for each pixel. Could each section correspond to a palette used by Rembrandt? This will be further investigated.

Additionally when taking a look at the sections 12 – 24 in figure 17 it can be seen that many of these sections also very localized in their original positions. These are likely the result of small changes or reparations over the lifetime of the artwork. It should be mentioned that elements that are typically associated with restoration such as Zn and Ti were *not* considered by t-SNE. Many of these changes are also found where the cap is painted. This suggests that the cap was modified in the lifetime of the artwork. Also these sections will be discussed in detail later on in this work.

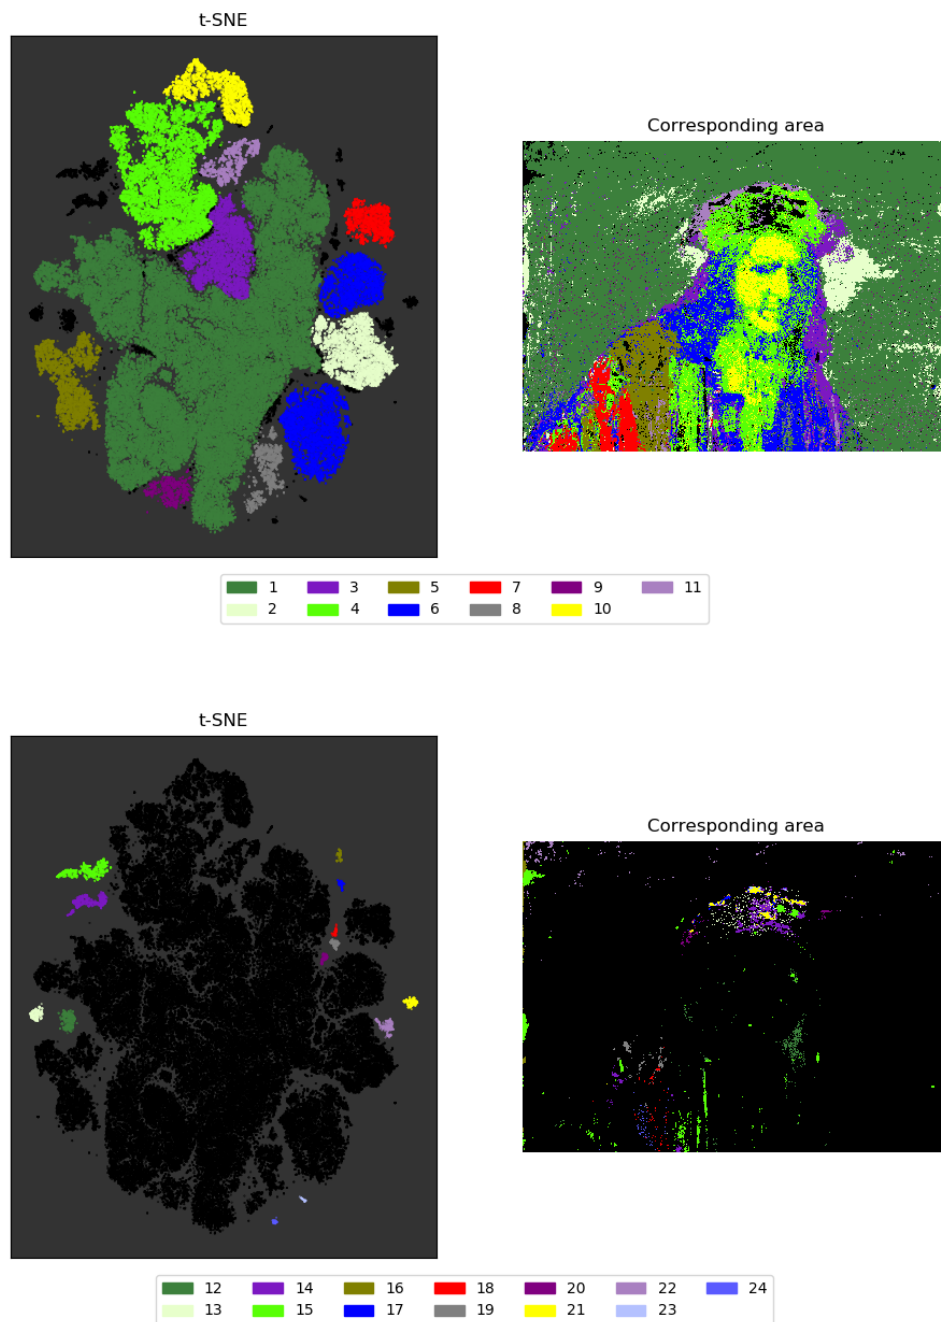


Figure 17: false colour representation of the various numbered t-SNE clusters mapped to their original position on the painting. The major figurative parts of the painting are in the large clusters from 1-11 are still easily recognized. The smaller and more localized clusters 12-24 depict the areas in which retouches are suggested by t-SNE.

Comparison

In order to verify whether the sections were chosen correctly, further analysis is needed. As mentioned earlier a section represents a number of pixels containing 19 different spectral peak intensities. For a number of pixels an average and standard deviation of each of the 19 spectral peak intensities can be calculated. This average ($\mu_{i,u}$) and standard deviation ($S_{i,u}$) also mentioned in equation 13 has is calculated for all the 24 sections. Subsequently for each section the mean spectral peak intensities have been compared by compared by statistical means.

Using a student *t*-test it can be easily concluded that, as expected, the average compositions of all sections is different from one another due to a large number of pixels per section. However it is more

interesting whether sections do significantly differ in all their 19 elemental intensity ratios, and for this the χ^2 -test is a more suited tool. Additionally it can be determined which elements give rise to the largest contributions to the $\hat{\chi}^2$ -statistic and thus which elements define the difference between two sections most prominently.

The χ^2 -test has been executed for every possible pair of sections, using a critical value of $p < 0.05$. The results of this test are summarized in table 4. The full table including p -values can be found in appendix C.

A few points should be noted when looking at the table. It is clear that not every combination of section is different when according to the χ^2 -test. The main reasons for this behaviour can be explained by the fact that within each section there still is a rather large standard deviation. This large deviation can be explained for the sections 1 – 11 since these are quite large sections. But for the sections 11 – 24, this is not the case and difference needs to be investigated further.

Table 4: Overview of results of the χ^2 -test on the intensity ratio differences between all sections. The cells marked with a “✓” are the all pairs of sections where with critical value $p < 0.05$. For these cells it is concluded that the compared sections are different. For an extended table with numbers see appendix C.

Section	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
1				✓	✓		✓			✓		✓	✓	✓	✓		✓	✓			✓	✓	✓	✓
2				✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓		✓	✓	✓		✓		✓	✓
3									✓	✓			✓					✓				✓		✓
4	✓	✓						✓	✓								✓	✓	✓	✓		✓		✓
5	✓	✓						✓	✓	✓										✓		✓	✓	
6		✓					✓			✓			✓	✓				✓				✓		✓
7	✓	✓				✓		✓	✓	✓		✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓
8				✓	✓		✓			✓		✓	✓	✓	✓		✓	✓				✓	✓	✓
9		✓	✓	✓	✓		✓			✓	✓	✓	✓	✓	✓		✓	✓	✓		✓	✓	✓	✓
10	✓	✓	✓		✓	✓	✓	✓	✓		✓	✓				✓	✓	✓	✓	✓	✓	✓	✓	✓
11		✓							✓	✓													✓	
12	✓	✓					✓	✓	✓	✓			✓	✓			✓	✓	✓	✓	✓	✓	✓	✓
13	✓	✓	✓			✓	✓	✓	✓			✓				✓	✓	✓	✓	✓	✓	✓	✓	✓
14	✓	✓				✓	✓	✓	✓			✓				✓		✓	✓	✓		✓		✓
15	✓	✓						✓	✓											✓		✓		
16							✓			✓			✓	✓				✓				✓		✓
17	✓	✓		✓			✓	✓	✓	✓		✓	✓					✓		✓		✓		✓
18	✓	✓	✓	✓		✓	✓	✓	✓	✓		✓	✓	✓		✓	✓			✓	✓	✓	✓	✓
19		✓		✓			✓		✓	✓		✓	✓	✓						✓		✓	✓	✓
20				✓	✓		✓			✓		✓	✓	✓	✓		✓	✓	✓		✓	✓	✓	✓
21	✓	✓					✓		✓	✓		✓	✓					✓		✓		✓		✓
22	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
23	✓	✓			✓		✓	✓	✓	✓		✓						✓	✓	✓		✓		✓
24	✓	✓	✓	✓		✓	✓	✓	✓	✓		✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓

Differences between t-SNE and chi-squared testing

What is can be perceived as a different group depends whether you are using the χ^2 -test or t-SNE. The reason behind this has a lot to do with the complexities when dealing with high dimensional data. An example where χ^2 -test and t-SNE would give different results is shown figure 18 [24]. In this case two interlinked circles are represented properly by t-SNE as two separate entities. However when a

χ^2 -test is performed between the x, y and z positions of the blue points their and the position and orange points, the χ^2 -test will not result in a significant difference. One can imagine that t-SNE tries to pull out all strings of points from the high dimensional space and laying them separately down on a 2d surface. If this is done in figure 18, and since both circles do not touch each other they can be pulled out separately and placed as two individual clusters in the t-SNE image.

The execution of a χ^2 -test between the orange and blue points can imagined as a test whether the centre of mass of the blue points and the orange points are far apart. The exact distance they have to be apart depends on the critical-value chosen and the standard deviation of the points in each axis (in the example a measure for the size of the circle).

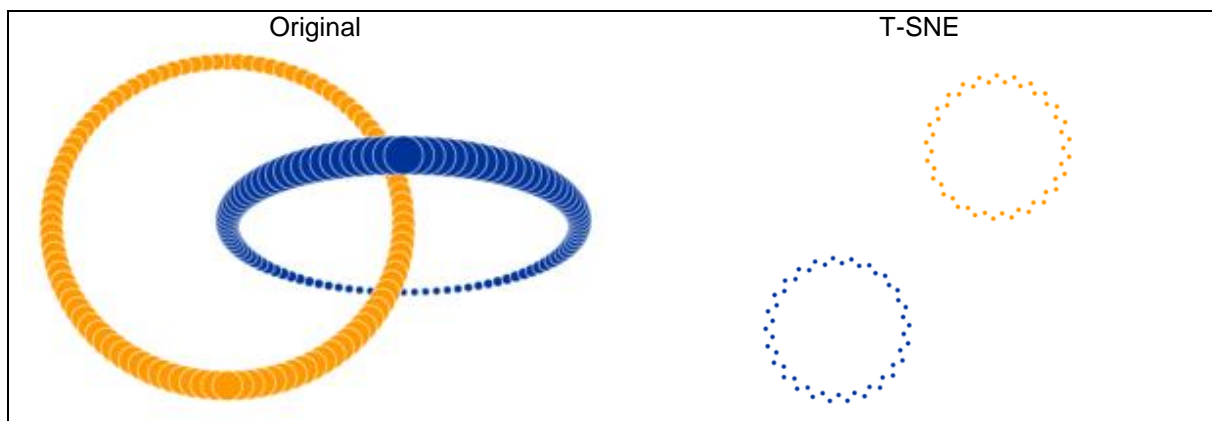


Figure 18: Showing the t-SNE representation of two linked circles. Although the centres of mass of both linked circles are not significantly far away enough for a chi-squared test to be considered a separate distribution. In the t-SNE representation both circles are nicely identified and visualized [24].

In the 19 dimensions for the higher dimensions it is not unlikely that there are strings of points that do not touch each other, yet have a similar centre of mass. If this is the case for two sections they will show up separate in the t-SNE image, but are not statistically significant different according to the χ^2 -test.

When two sections are different according to both t-SNE clustering and the χ^2 -test, then it can safely assumed that these should be considered separate. However when the according to t-SNE the section is different but not according the χ^2 -test, caution is needed. However this does not necessarily mean that t-SNE has given a false positive. These groups might still be different, but this must be based upon other evidence. Such evidence is for example that the two sections are depicting something else entirely (and even have a different colour) in the painting.

A different method of quantifying differences between sections might prove useful. For example it could be useful to use a graph that represents each point as a node and the connections of the nodes are based on a certain similarity between data points. In t-SNE already such a graph is constructed [4], that could be used for this purpose. Within this graph from each point a random walk is performed with K steps. The chance that the random walk in a node is recorded for each node. It is hypothesized that a random walk has a reasonable chance to end up within the same section from which it started. This chance that a random walk from a point ends up in a different section could be used as a probability in a similar fashion. This method can be done by doing explicitly doing random walks. However the PageRank algorithm can also be adapted in order to achieve significant speed benefits [25].

Finally it should be kept in mind that a full paint composition cannot be derived from XRF scanning. The presence or lack of organic pigments might be the reason that a painting mixture is different that both t-SNE and the χ^2 -test will fail to recognize. In this research we deem it unlikely that this ever was the case. However it is possible to add in more data from different imaging techniques such as infra-red or a photograph to enhance separation by t-SNE.

Sections within the painting

With the painting divided into section the most interesting questions remain unanswered: do these sections show palettes? And what can we learn about the painting from these sections? This last section of the report will discuss the meaning of these sections and show if indeed palettes are found.

Dead paint and background

One of the more interesting comparisons between the section to make is a comparison between section 1 and section 2. The major difference between the sections is that section 2 has a higher concentration of elements (Co, Ni) that can be associated with smalt. When looking at the visual painting the colour however with the sections is nearly identical compared with those of the background. Also when looking at the shape of the patch of section 2 this suggests that the smalt is present in an underpainting. This is likely to create a shadow around the head to create a sense depth. So that the head of the man appears separated from the background. Interestingly enough the same composition can also be found on the righthand side of the painting. Which suggests that also on these parts some smalt in the underpainting could be present. Or that some horizon is attempted to be constructed.

In another work of Rembrandt, *An painting of an Old man* a similar shadow can be seen in the dead paint [26]. For example when the dead paint layer as visualized by the Cu elemental map as seen figure 22 is compared with section 2, it can be seen that these two are quite similar in what they represent.




Figure 19: An Cu elemental distribution map of *Portrait of an Old Man*, Rembrandt van Rijn, ca. 1630, oil on wood, private collection [26].

Also current cross-sections taken from this painting seem to support the idea that there is a smalt layer present in the underpainting. Whereas the background layers are then painted on top of that. In figure 20 a cross section is shown where there's still smalt in the background [27].




Figure 20: Dark field light microscopy image of a cross section taken from an area around the hat. The position where the sample was taken is shown on the right side. In the cross section 4 layers can be distinguished: these are (0) a quartz ground, (1) a layer of discoloured smalt, (2) a dark organic paint layer and (3) a overpainted layer rich in vermillion [27].

Background to foreground

When taking a look at section 3 (see also figure 21), it can be seen that most of these mostly depict the contour of the man in the painting. The majority of section 3 does not overlap with the background, but seems to present at the inside of the man.



Figure 21: Section 3 highlighted in green over the visual image of the man with the red baret. Most of section 3 sits at the edge of the background and the foreground.

One could wonder whether this would be a separate palette or whether this section only exists because the background and foreground overlap in section 3. It should be remembered after all that XRF-radiation can penetrate into deeper layers that might not be visible. A way to find out is to take a look at the recorded intensities of some major spectral lines. In table 1 some average spectral intensities are listed. What can be then is that for some major elements such as Co (smalt), Fe (earth pigments) and Pb (lead-white) the average intensity of sections 3 sits right in between the sections 1 (background) and section 4 (clothing).

Table 5: Average spectral line intensities of three major elements for sections 1, 3, and 4.

Section	Co K α	Fe K α	Pb L β
Section 1	5.51	28.7	3.19
Section 3	4.04	17.4	9.94
Section 4	0.98	13.2	31.5

This average suggests that section 3 is some mixture of the sections 1 and 4. The only question remains whether this is part of the man stacked upon the background or whether the paint is truly mixed. The stacking of layers seems more plausible, for a number of reasons. First of all in the visual image this section cannot be distinguished. If the paint were truly mixed, it is expected that this section would also have a different colour than its surroundings. Second it is known that Rembrandt usually began painting the background of the scene and would go from there to the foreground of the scene. And if the background and the foreground are of a different palette it would make sense that a part of the foreground can be painted over the background. This would also explain why section 3 is also fully in the inside of the man. Thus it is concluded section 3 only exists because here a foreground layer is painted over a background layer. This could be reinforced by taking a cross-section in one of the section 3 regions.

Face and cloth

It's also expected that some scenes in the foreground need a change of palette. It's obvious that the face of the man needs a different paint mixture than the clothes of the man. From the sections 5, 7 and 11 it can be derived that also clothes of the man were painted with several different pigment compositions. However in contrast to what is seen with section 3, there seems to be no overlap between the sections 5, 7 and 11. It's not fully clear as to why there is overlap between the background and foreground and not between the clothing sections of the man. An explanation could be that these clothing sections were painted simultaneously and the painter deemed it undesirable that paint between these sections would mix while the paint was still drying. However what *is* clear is that the man is made up of a limited number of sections, which can be very well explained by that a limited number of paint mixtures was used to paint the man.

What about the red cap?

At first glance there are many different spectral compositions present in the cap. The main reason is that the cap has been overpainted. This can also be seen when comparing the Pb and the Fe elemental distributions as shown in figure 22. However there are also a lot of very localized unique paint mixtures from the sections 12-24 present in the cap. This also suggests that many smaller changes have been made to the cap. For example visually the cap and the cloth on the man on the left side have roughly the same colour red. However the elemental composition is very different. This suggests that the cap and the cloth were painted separately. Finally in a cross-section an intermediate layer of varnish was found further supporting that the cap was overpainted [27]. This varnish layer can be seen in figure 23.




Figure 22: Fe K α and Pb L β elemental spectra. These elemental distributions show a rather different shape of hat on the man.



Figure 23: Dark field light microscopy image of a cross section taken from an area around the hat. The position where the sample was taken is shown on the right side. In the cross section 9 layers can be distinguished: these are (1) yellow layer with partially soaped lead-tin-yellow (2), a layer with lead-white (3) dark irregular layer with bone black, smalt, yellow and red ochre and fine red earth. (4) Thick, compact layer of red ochre and bone black. (5, 6) Old varnish layers. (7, 8) Overpainted red layers with varying concentration of vermilion and (9) The varnish layer [27].

That the cap is overpainted is clear. However the large number of sections make it unclear whether this is a single overpainting or whether this cap has been overpainted multiple times. Or even overpainted and then restored. Typical elements that are associated with restoration are Ca, Zn and Ti. It should be remembered that both Zn and Ti were excluded for t-SNE, but not Ca which is also associated with restorations. If all the pixels where Ca was present are discarded the changes that were due to very early restorations or modifications can be shown. In figure 24 this is shown. It can be seen that indeed not only the cap has seen changes, but also the coat of the man on the left. This gives an indication as to which regions of the painting have been altered, besides restorations. This might also prove to give more insight into changes in the painting. What however cannot be learned from purely technical means is who made this modification and why.




Figure 24: t-SNE sections 12 - 24. Where no Ca, Ti or Zn (pigments associated with restorations) are present.

Conclusions

The aim of this work was to investigate and identify the use of palettes in the *Man with the red baret*. By determining the various paint mixtures that are found spatially over the painting it can be seen that indeed a limited number of paint mixtures was used to construct the work of art. In fact the majority of the painting can be reduced to 12 separate sections. Furthermore when it is considered that some sections are there due to overlap between background and foreground, the number of steps required to create painting is more likely to be around 10. The small number of steps appears to confirm the hypothesis that indeed a very systematic approach was used for this painting.

Additionally in this work it is shown that, while computationally intensive, it is feasible and possible to perform fitting for every element in a XRF spectrum and that the result is internally consistent. Also the fitness parameters (χ^2 and ψ^2) seem to match what is statistically expected. Model fitting the XRF spectra seems to make it possible to visualize otherwise invisible elements such as sulphur. Secondly, it helps filtering out background radiation in order to do proper statistical analysis later on. It should also be noted that background radiation can be quite significant and is not uniformly distributed (figure 5) over the painting as well, and thus can also create a distorted view when looking at the elemental distribution maps. Even the elemental maps as a result of fitting are to be preferred over unfitted elemental maps.

Visualising the spectral lines using t-SNE is a useful aid in recognizing regions of interest in the painting. However a painting cannot be divided into separate regions/sections by using t-SNE alone. Also figurative units have to be taken into account. The statistically significant difference between two painting sections can be calculated using the chi-squared test. This proves indeed that many clusters are unique, but not all. This difference can be partially explained by considering what the chi-squared test does and what t-SNE does.

The existence of the regions can be very well explained when it is assumed that indeed palettes were used in the 17th century painting process. It allows to hypothesize in what order the various figurative units were painted and also give an excellent indication of what cross sections might look like. This is also affirmed by the cross-sections already taken.

Local and unique sections in a painting do suggest the presence of retouches. It was already possible to find many of these by looking at elements that are typically used for restoration such as Zn and Ti. However with the help of t-SNE also anomalies can be detected that do not contain any of these elements. This can assist in finding changes in the painting that have been made earlier in its lifespan.

Finally doing an in-depth analysis of XRF data can give much valuable insight in the working process of the paint. If anything, this study does show that doing XRF scanning is indeed a very powerful method for investigating paintings, and moreover that there is indeed much to be gained from looking into spectral compositions for each point in a painting.

Recommendations

In the discussion a few cross sections were shown and reinforce what is learned from the XRF data. However with the knowledge of where sections are more interesting regions can be found for taking a cross section. For example in section 3 was a result of paint layer stacking can be verified by taking a cross section. This verification in turn could corroborate the current working methodology.

The combining of model fitting the raw XRF data with t-SNE allows for new insights in the analysis of paintings. One of the nice properties of t-SNE is that it allows more than just elemental distributions to be considered as high dimensional input. For example a photograph containing the visual data of the painting and hyperspectral data can also be used as additional inputs to effectively and more reliably analyse the painting. Also the chi-squared test is dimensionless and allows for multiple data sources to be used simultaneously.

An alternative to the chi-squared test could be used in order to quantify the difference between t-SNE clusters. Such a method could involve further analysing the underlying graph that t-SNE also uses when calculating the Kullback-Leibler divergence. In this graph a random walk with K steps from a node is made and for each node the chance that the walk ends there, can then be calculated. Using this method the chance that a point belongs to a cluster can then be quantified.

Most of the data analysis was done on the top part of the painting. However it should not be forgotten that the *man with the red baret* also has a bottom part and also that part of the painting should be considered when recreating the construction steps of this painting. In order to perform t-SNE successfully a method must be devised for stitching either raw data or fitted spectra.

Finally this analysis of fitting, t-SNE and comparing sections to one another has only been performed on a single artwork for this work. All techniques used are directly applicable on other artworks as well. The methods used in this work can also be applied to other works of Rembrandt as well to further identify the use of palettes by Rembrandt. But of course this should not be restricted to purely Rembrandt or painters from the time period, it might also be interesting how paint mixtures vary for paintings from other time periods. All in all, using these methodologies might uncover a whole new level of insight in the paintings of the 17th century.

References

- [1] The Rembrandt Database, „History of conservation & research, attributed to Rembrandt, or studio of Rembrandt, Man with red baret, early 1660s, Museum Boijmans Van Beuningen, Rotterdam, inv./cat. 2113 (OK),” [Online]. Available: <http://www.rembrandtdatabase.org/Rembrandt/painting/2936/man-with-red-baret>. [Geopend 16 Februari 2018].
- [2] E. v. d. Wetering, „Reflections on the relation between technique and style : the use of the palette by the seventeenth-century painter,” in *Historical Painting Techniques, Materials, and Studio Practice. Preprints of a Symposium*, 1995.
- [3] M. Mantler en M. Schreiner, „X-ray fluorescence spectrometry in art and archaeology,” *X-Ray Spectrometry*, vol. 29, pp. 3-17, 1 2000.
- [4] L. J. P. van der Maaten en G. E. Hinton, „Visualizing High-Dimensional Data Using t-SNE,” *Journal of Machine Learning Research*, vol. 9, pp. 2579-2605, 2008.
- [5] W. Beurs, „De Groote Waereld in het Kleen Geschildert,” Amsterdam, , 1698, pp. 165-166.
- [6] K. Groen, „Investigation of the use of the binding medium by Rembrandt: chemical analysis and rheology,” *Zeitschrift für Kunsttechnologie und Konservierung*, 1997.
- [7] R. J. Gettens en G. L. Stout, „Painting Materials: A Short Encyclopedia,” Dover Publications, 2012, p. 319.
- [8] D. J. Kalnicky en R. Singhvi, „Field portable XRF analysis of environmental samples,” *Journal of Hazardous Materials*, vol. 83, pp. 93-122, 5 2001.
- [9] A. C. Thompson, D. T. Attwood, E. M. Gullikson, M. R. Howells, J. B. Kortright, A. L. Robinson, J. H. Underwood, K.-J. Kim, J. Kirz, I. Lindau, P. Pianetta, H. Winick, G. P. Williams en J. H. Scofield, *X-Ray Data Booklet*, A. C. Thompson, Red., Lawre, 2009.
- [10] J. Inczedy, T. Lengyel, A. M. Ure, A. Gelencsér en A. Hulanicki, „Compendium of analytical nomenclature,” *The Orange Book, 3rd Edn.*, 1998.
- [11] J. R. Sieber, „X-rays in research and development at the National Institute of Standards and Technology,” *X-Ray Spectrometry*, vol. 29, pp. 327-338, 2000.
- [12] F. A. Gianturco en C. A. Coulson, „Inner-electron binding energy and chemical bonding in sulphur,” *Molecular Physics*, vol. 14, pp. 223-232, 1 1968.
- [13] W. S. J. Taft en J. W. Mayer, *The Science of Paintings*, Springer, 2013.
- [14] D. Bomford, „Rembrandt,” National Gallery, 2006, pp. 35-47.
- [15] H. Kühn en H. Kuhn, „Trace Elements in White Lead and Their Determination by Emission Spectrum and Neutron Activation Analysis,” *Studies in Conservation*, vol. 11, p. 163, 11 1966.
- [16] R. De la Rie, „Old Master Paintings: A Study of the Varnish Problem,” *Analytical Chemistry*, vol. 61, pp. 1228A--1240A, 11 1989.
- [17] P. Dietemann, C. Higgitt, M. Kälin, M. J. Edelmann, R. Knochenmuss en R. Zenobi, „Aging and yellowing of triterpenoid resin varnishes – Influence of aging conditions and resin composition,” *Journal of Cultural Heritage*, vol. 10, pp. 30-40, 1 2009.

- [18] M. Alfeld, J. V. Pedroso, M. Eikema Hommes, G. V. Snickt, G. Tauber, J. Blaas, M. Haschke, K. Erler, J. Dik en K. Janssens, „A mobile instrument for in situ scanning macro-XRF investigation of historical paintings,” *Journal of Analytical Atomic Spectrometry*, vol. 28, p. 760, 2013.
- [19] Kaye & Laby Online, „4.2.1 X-ray absorption edges, characteristic X-ray lines and fluorescence yields,” in *Tables of Physical & Chemical Constants*, The National Physical Laboratory, 2005.
- [20] W. Press, „Numerical recipes : the art of scientific computing,” Cambridge, UK New York, Cambridge University Press, 2007, pp. 730-734.
- [21] B. J. Thijsse, M. A. Hollanders en J. Hendrikse, „A practical algorithm for least-squares spline approximation of data containing noise,” *Computers in Physics*, vol. 12, p. 393, 1998.
- [22] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot en E. Duchesnay, „Scikit-learn: Machine Learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825-2830, 2011.
- [23] L. Van Der Maaten, „Accelerating t-SNE using tree-based algorithms.,” *Journal of machine learning research*, vol. 15, pp. 3221-3245, 2014.
- [24] M. Wattenberg, F. Viégas en I. Johnson, „How to Use t-SNE Effectively,” *Distill*, 2016.
- [25] A. Rajaraman en J. D. Ullman, „Link Analysis,” in *Mining of Massive Datasets*, Cambridge University Press, pp. 139-175.
- [26] M. Alfeld, D. P. Siddons, K. Janssens, J. Dik, A. Woll, R. Kirkham en E. Wetering, „Visualizing the 17th century underpainting in Portrait of an Old Man by Rembrandt van Rijn using synchrotron-based scanning macro-XRF,” *Applied Physics A*, vol. 111, pp. 157-164, 01 4 2013.
- [27] A. Loon en K. Keune, „Analytisch chemisch verfmonsteronderzoek Rembrandt Man met rode muts Museum Boijmans van Beuningen,” 2007 .
- [28] V. A. Solé, „The PyMca Application and Toolkit,” in *Python for Scientific Computing Conference*, 2009.
- [29] V. A. Solé, E. Papillon, M. Cotte, P. Walter en J. Susini, „A multiplatform code for the analysis of energy-dispersive X-ray fluorescence spectra,” *Spectrochimica Acta Part B: Atomic Spectroscopy*, vol. 62, pp. 63-68, 1 2007.

Appendix A: Extended table of fitting parameters

The fitting parameters used for fitting each pixel. The parameters are found from model fitting the sum spectrum. The model used is described in equation 2: $C(x) = a + bx + H_1 e^{-\frac{1}{2} \left(\frac{x-H_1}{w_1} \right)^2} + H_2 e^{-\frac{1}{2} \left(\frac{x-H_2}{w_2} \right)^2} + \dots$ In order to come to the initial guess for each individual pixel the parameters a , b and H_n need to be divided by the number of pixels to be fitted. This means that both the binned and the unbinned pixels can both be fitted using this table.

Energy Range	Channels	Peak	μ	Constraints	w [stdev]	Constraints	Height	Constraints	a	Constraints	b	Constraints	k	v	
2.1243 - 2.5643	308 - 351	Ska	2.308	Fixed	0.0539	0.045 - 0.055	189885		>0					11	
		PbMa	2.345	Fixed	0.0464	0.045 - 0.055	398251		>0						
3.1143 - 3.4743	407 - 442	Ka	3.325	Fixed	0.0574	0.04 - 0.07	428200		>0	2.83E+04	None	8.64E+04	None	33	
		ArKb	3.192	Fixed	0.0457	0.04 - 0.07	141721		>0	1.21E+06	None	-2.64E+05	None		
3.5043 - 3.8243	446 - 477	CaKa	3.699	3.5043 - 3.8243	0.0597	0.04 - 0.07	6.06E+06		>0	4.40E+06	None	-1.13E+06	None	5	
5.7043 - 6.0343	666 - 698	MnKa	5.900	Fixed	0.0612	0.04 - 0.07	365286		>0	-1.96E+06	None	4.50E+05	None	4	
6.0443 - 6.5743	700 - 752	Feka	6.400	Fixed	0.0629	0.04 - 0.07	5.15E+07		>0					6	
		MnKb	6.491	Fixed	0.0400	0.04 - 0.07	4.03E+06		>0	-7.48E+06	None	1.37E+06	None		
6.6743 - 7.284	763 - 823	Coka	6.925	Fixed	0.0703	0.03 - 0.10	1.00E+06		>0					6	
		Fekb	7.058	Fixed	0.0659	0.03 - 0.10	7.92E+06		>0	1.00E+06	None	-2.72E+03	None		
7.2943 - 7.8443	825 - 879	NiKa	7.480	Fixed	0.0692	0.045 - 0.08	1.42E+06		>0					6	
		Cokb	7.649	Fixed	0.0689	0.045 - 0.08	1.24E+06		>0	-3.85E+04	None	1.35E+05	None		
7.8543 - 8.1743	881 - 912	CuKa	8.045	Fixed	0.0639	0.04 - 0.07	3.64E+05		>0	-5.74E+06	None	8.62E+05	None	4	
		HgLa1	9.999	Fixed	0.0671	0.04 - 0.07	4.88E+05		>0					4	
9.7043 - 10.154	1066 - 1110	HgLa2	9.899	Fixed	0.0465	0.04 - 0.07	9.80E+04		>0	-3.25E+06	None	4.92E+05	None	6	
10.164 - 10.974	1112 - 1192	PbLa/AsK	10.543	Fixed	0.0824	0.03 - 0.10	2.70E+07		>0	2.55E+06	None	-5.38E+04	None	4	
11.494 - 12.004	1245 - 1295	HgB	11.823	Fixed	0.0991	0.03 - 0.10	3.61E+05		>0					6	
		AsKb	11.726	Fixed	0.0799	0.03 - 0.10	2.25E+06		>0	9.06E+05	None	7.85E+04	None		
12.334 - 12.874	1329 - 1382	PbLb	12.617	Fixed	0.0824	0.03 - 0.10	1.68E+07		>0	-1.11E+07	None	1.09E+06	None	4	
14.504 - 14.974	1546 - 1592	PbLg	14.762	Fixed	0.0800	0.03 - 0.10	1.74E+06		>0	-2.51E+06	None	2.70E+05	None	4	
24.004 - 25.984	2496 - 2693	SnKa	25.175		25.0 - 25.3	0.1714	0.10 - 0.20	3.57E+04		>0	6.58E+05	None	-2.19E+04	None	5

193

Appendix B: Elemental distribution maps

This appendix shows the elemental map for each element. All the images are a result of model fitting a spectral intensity peak. Hence the background radiation component is removed. A pixel that is white has a high spectral intensity value. Whereas that a pixel that is black has low intensity value. Note that the contrast for all images is *maximized*, this means that black does not have to correspond with an intensity of 0. Additionally it means that the same brightness in one image has a different intensity associated with it in another image.

As K β

Top



Bottom



Ca K α

Top



Bottom



Co K α

Top



Bottom



Ca K β

Top



Bottom




Cu Ka

Top




Bottom



Fe Ka

Top



Bottom



Fe K β

Top



Bottom



Hg L α

Top




Bottom



Fe L β

Top



Bottom




K K α

Top



Bottom




Mn K α

Top



Bottom



Ni K α

Top



Bottom



Pb L α & As K α

Top



Bottom



Pb L β

Top



Bottom



Pb L γ

Top



Bottom



Pb Ma

Top



Bottom



Pb Mβ

Top



Bottom




S Kα

Top



Bottom



Sn K α

Top



Bottom




Zn K α

Top



Bottom



Appendix C: extended χ^2 -test result values

Overview of χ^2 -test right-tail probabilities. The number in the table represents the likelihood that two sections are drawn from the same distribution. Values < 5% are marked and the corresponding sections and are considered to be unrelated to one another.

Section	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
1		0,611	0,214	0,000	0,000	0,744	0,000	1,000	0,999	0,000	0,090	0,000	0,000	0,000	0,008	0,999	0,005	0,000	0,084	1,000	0,024	0,022	0,000	0,000
2	0,611		0,505	0,000	0,000	0,003	0,001	0,337	0,000	0,000	0,004	0,000	0,000	0,000	0,187	0,000	0,000	0,000	0,000	0,325	0,000	1,000	0,000	0,000
3	0,214	0,505		0,361	0,401	0,987	0,181	0,079	0,001	0,000	0,798	0,021	0,068	0,896	0,945	0,175	0,002	0,000	0,107	0,361	0,197	0,003	0,245	0,000
4	0,000	0,000	0,361		0,934	0,121	0,074	0,000	0,939	0,900	0,275	0,095	0,099	1,000	0,175	0,001	0,002	0,006	0,006	0,000	0,080	0,000	0,059	0,000
5	0,000	0,000	0,401	0,934		0,139	0,158	0,000	0,006	0,996	0,054	0,063	0,035	0,987	0,399	0,990	0,035	0,210	0,934	0,000	0,192	0,000	0,46	0,980
6	0,744	0,003	0,987	0,121	0,139		0,000	0,796	0,356	0,651	0,833	0,004	0,035	0,647	0,990	0,740	0,004	0,000	0,000	0,460	0,540	0,000	0,305	0,000
7	0,000	0,001	0,181	0,074	0,158	0,000		0,000	0,000	0,775	0,000	0,001	0,001	0,544	0,740	0,618	0,000	0,000	0,000	0,982	0,997	0,000	0,305	0,000
8	1,000	0,337	0,079	0,000	0,000	0,796	0,000		1,000	0,000	0,065	0,000	0,000	0,003	1,000	0,012	0,000	0,000	0,201	0,057	0,003	0,000	0,000	0,000
9	0,999	0,000	0,001	0,000	0,000	0,356	0,000	1,000		0,000	0,009	0,000	0,000	0,000	0,999	0,734	0,000	0,000	0,000	0,894	0,004	0,000	0,000	0,000
10	0,000	0,000	0,000	0,000	0,000	0,006	0,000	0,000		0,046	0,000	0,251	0,401	0,734	0,000	0,740	0,000	0,000	0,000	0,000	0,000	0,000	0,997	0,000
11	0,090	0,004	0,919	0,900	0,996	0,651	0,775	0,065	0,009	0,046	0,185	0,185	0,940	0,927	0,991	0,740	0,991	0,994	0,948	0,159	0,907	0,000	0,944	0,577
12	0,000	0,000	0,798	0,275	0,054	0,833	0,000	0,000	0,000	0,185		0,004	0,012	0,971	0,618	0,000	0,000	0,000	0,000	0,460	0,540	0,000	0,305	0,000
13	0,000	0,000	0,021	0,995	0,063	0,004	0,001	0,000	0,000	0,251	0,004		1,000	0,992	0,000	0,026	0,000	0,000	0,000	0,000	0,000	0,000	0,081	0,000
14	0,000	0,000	0,068	0,997	0,092	0,035	0,001	0,000	0,000	0,401	0,927	0,012	1,000		0,991	0,154	0,000	0,000	0,000	0,000	0,000	0,000	0,540	0,000
15	0,008	0,000	0,896	1,000	0,987	0,647	0,544	0,003	0,000	0,734	0,991	0,971	0,992	0,991		0,812	0,300	0,111	0,198	0,012	0,613	0,000	0,860	0,167
16	0,999	0,187	0,945	0,175	0,399	0,990	0,035	1,000	0,999	0,000	0,740	0,618	0,000	0,001	0,812		0,246	0,000	0,294	0,982	0,456	0,074	0,000	
17	0,005	0,000	0,148	0,002	0,080	0,241	0,000	0,012	0,000	0,000	0,000	0,026	0,154	0,300	0,246		0,000	0,276	0,744	0,000	0,997	0,768	0,000	
18	0,000	0,000	0,000	0,002	0,961	0,000	0,002	0,000	0,000	0,994	0,000	0,000	0,000	0,111	0,000	0,000	0,000		0,000	0,000	0,000	0,000	0,018	
19	0,084	0,000	0,107	0,006	0,491	0,210	0,000	0,201	0,012	0,948	0,000	0,000	0,000	0,198	0,294	0,744	0,276		0,001	0,065	0,000	0,000	0,000	
20	1,000	0,325	0,792	0,000	0,000	0,460	0,000	1,000	0,894	0,000	0,159	0,000	0,000	0,012	0,982	0,000	0,000	0,001		0,000	0,000	0,000	0,000	
21	0,024	0,000	0,197	0,080	0,192	0,540	0,000	0,057	0,004	0,000	0,907	0,001	0,045	0,170	0,613	0,456	0,997	0,000	0,065	0,000		0,997	0,000	
22	0,022	1,000	0,003	0,000	0,000	0,000	0,000	0,003	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000		0,000	0,000	
23	0,000	0,000	0,245	0,059	0,046	0,305	0,000	0,000	0,000	0,000	0,944	0,003	0,081	0,540	0,860	0,074	0,768	0,000	0,000	0,997		0,000	0,000	
24	0,000	0,000	0,000	0,018	0,980	0,000	0,000	0,000	0,000	0,000	0,577	0,000	0,000	0,000	0,000	0,000	0,000	0,018	0,000	0,000	0,000	0,000	0,000	

Appendix D: extended report per section

In this appendix a small report is given in for each separate section.

In the top left an image is shown where the pixels are located of the section in the original painting, while in the top right side it is shown where the pixels are located after t-SNE transformation.



The *average intensity* shows two intensities. In blue it is the average spectral line intensity of all the points that belong to the section. The black bars is the spectral line intensity of all points, including the points for the section. This means that the height of the black bars is the same in each chart.

The *Relative average intensity* is the average intensity of a spectral line of all points in a section, divided by the average spectral line intensity for all points. E.g. If the average Fe K α intensity is 70 in a section, and the average Fe K α intensity is 40 for all pixels, then the *Relative average intensity* is 70/40 for Fe K α for section 1. This relative intensity is a simple method to compensate for the fact that some spectral lines are more intense than others. Alternatively this can be seen as the blue bar divided by the black bar in the *average intensity* plot.


The *Standard deviation* is the sample standard deviation of a section. In the blue bar the sample is the intensities of all the pixels that are part of the section. In the black bar the sample is every point in the painting. Thus for every *Standard deviation* chart the black bars have the same height.

The *Relative standard deviation* is the standard deviation of a spectral line of the points in a section divided by the standard deviation of a spectral line of all points. Alternatively one could see this as the blue bar divided by the black bar in the *Standard deviation* chart.


Section 1




Average intensity




Standard deviation





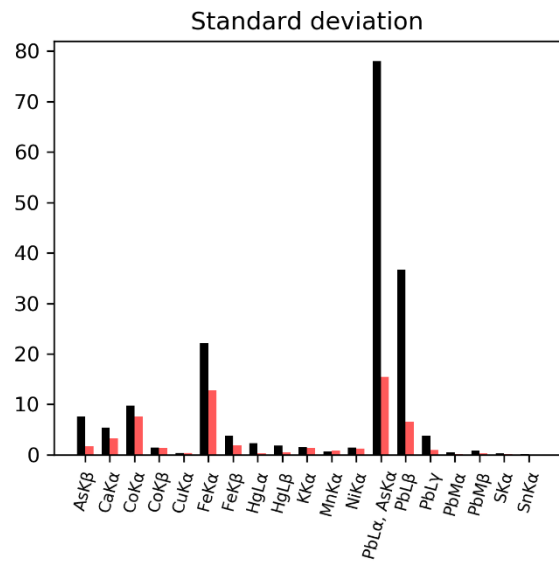



Relative average intensity









Relative standard deviation





Section 2




Section 3




Section 4




Average intensity




Standard deviation









Relative average intensity









Relative standard deviation



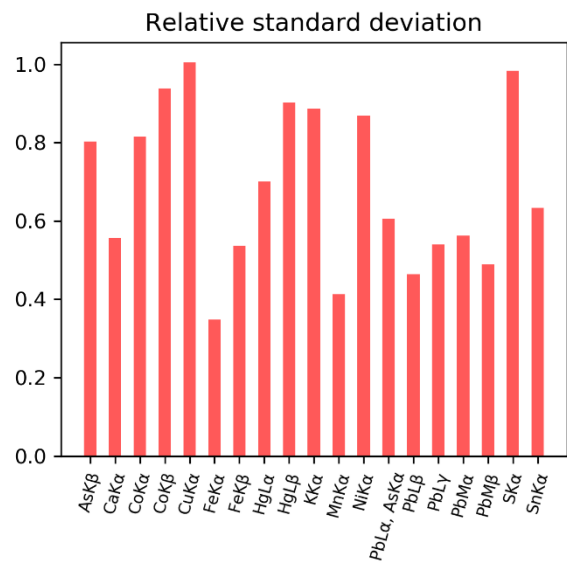





Section 5









Section 6





Section 7




Section 8




Section 9




Average intensity




Standard deviation




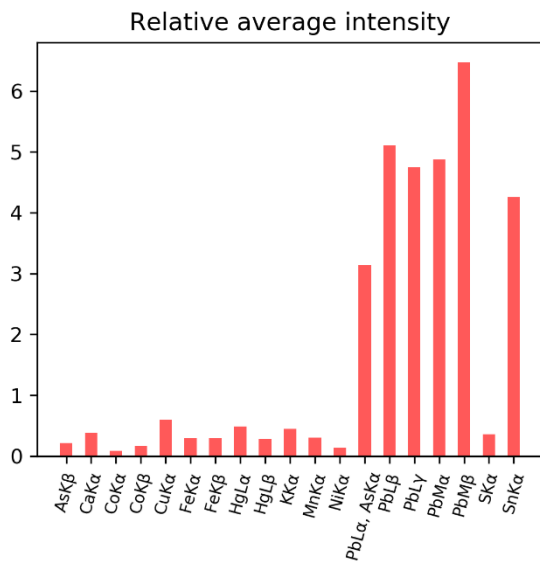



Relative average intensity





Relative standard deviation




Section 10




Section 11




Average intensity




Standard deviation









Relative average intensity









Relative standard deviation









Section 12





Section 13




Section 14




Section 15




Average intensity




Standard deviation








Relative average intensity









Relative standard deviation









Section 16









Section 17









Section 18




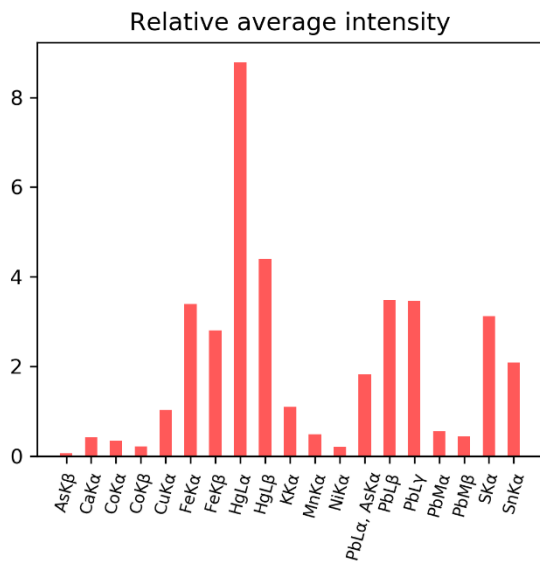




Section 19








Section 20








Section 21



Section 22



Section 23



Section 24

