VisuaLayered: Combined Visual Analysis of MA-XRF and RIS Data

Andra Popa

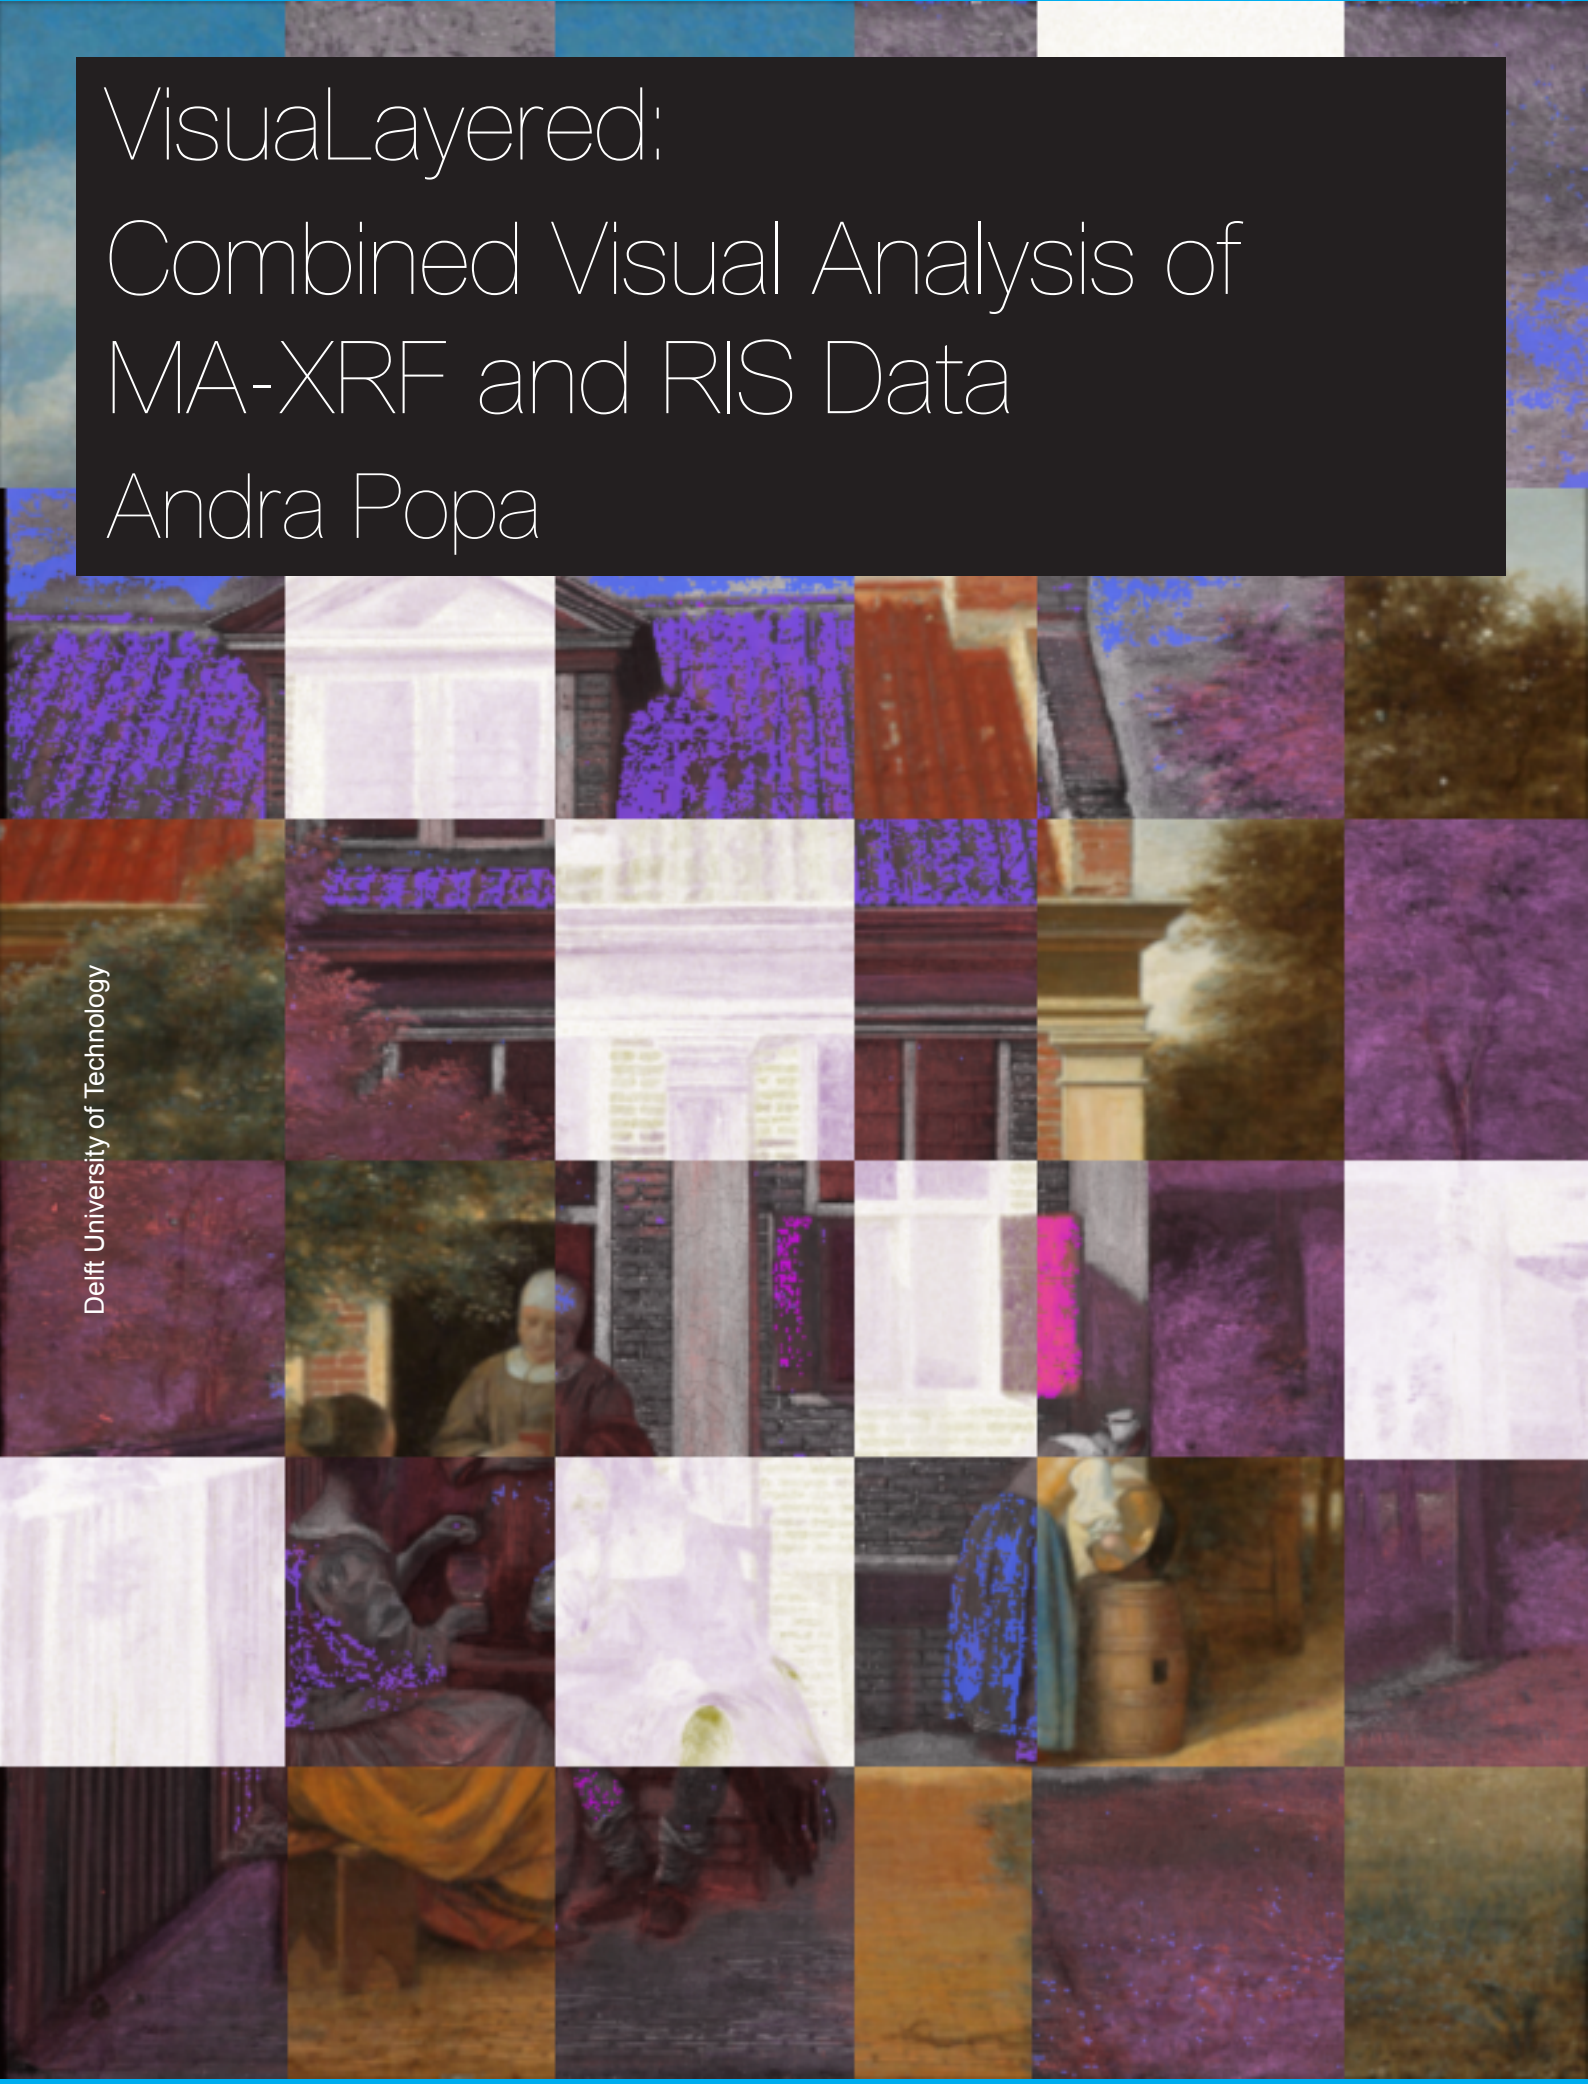

Delft University of Technology

TU Delft

VisuaLayered:

Combined Visual Analysis of MA-XRF and RIS Data

by

Andra Popa

In partial fulfilment of the requirements for the degree of **Master of Science**

in the Software Technology track

at the Faculty of Electrical Engineering, Mathematics and Computer Science of

Delft University of Technology,

to be defended publicly on Wednesday December 7, 2022 at 9:00 AM.

Student number:	5251249	
Project duration:	January 24, 2022 – December 7, 2022	
Thesis committee:	Prof. dr. Elmar Eisemann,	TU Delft, thesis advisor
	Catharine Oertel,	TU Delft
	Thomas Höllt,	TU Delft, daily supervisor

TU Delft

Abstract

In this work, we present VisuaLayered, the implementation of a combined analysis workflow for pigment identification. VisuaLayered is an integrated, interactive system that focuses on the combined visual analysis of Macro X-Ray Fluorescence (MA-XRF) and Reflectance Imaging Spectroscopy (RIS) data.

Analysing paintings for pigment identification is relevant for many applications in the cultural heritage domain, such as conservation and restoration. Domain experts use non-invasive scanning techniques as an initial step in their analysis. Two such techniques are MA-XRF and RIS. They provide hyperspectral data on the elemental and molecular composition of pigments, respectively. Domain experts analyse these two complementary data modalities in order to determine the pigments present in the different paint layers of a painting. However, due to the size and high-dimensionality of these datasets, experts have problems with efficiently analysing the data. In general, they examine the two data modalities separately in the analysis workflow and use their knowledge to unify all the information without additional software support, as there is no integrated system that is designed specifically for the combined analysis of MA-XRF and RIS data.

We worked in collaboration with domain experts from the Rijksmuseum, Amsterdam in order to design and implement VisuaLayered based on current domain practices. We use t-SNE for projecting the high-dimensional data into a two dimensional space, which the user can interactively explore in combination with other linked views in order to find connections between the two data modalities. With our system, experts can explore the spatial distribution and the correlation between pixels that have similar molecular and/or elemental compositions. Additionally, for the RIS data, we support endmember identification and analysis based on the pigments' spectral profiles. We tested the efficiency of our system with respect to the designed workflow in a case study evaluation with our collaborator. They successfully used VisuaLayered for the analysis of one painting and found the views combining the two data modalities very useful for better understanding the relation between them. Moreover, they were able to identify new pigments, that they missed when using existing software.

Preface

First of all, I would like to thank my supervisor, Thomas Höllt, for offering me the amazing opportunity of working on this project and for putting me in contact with our collaborators. Thank you for all the guidance offered at our weekly meetings, for the support and words of encouragement throughout the entire project that gave me more confidence in my work and for suggesting and helping me participate in two conferences, that turned out to be more fun and interesting than scary.

I would like to thank Francesca Gabrieli for sharing her expertise in cultural heritage with me, for her continuous feedback during the project, for her words of encouragement when I was not that confident in myself and for always being happy to answer any of my questions.

I would like to thank my thesis advisor, Elmar Eisemann, for his support and the Computer Graphics and Visualization group for their feedback on my presentations. Additionally, I would like to thank Thomas Kroes for taking the time to help me understand various aspects of HDPS.

I would also like to thank Matthias Alfeld for helping me understand the background of this work and for pointing out any mistakes I made when talking and writing about cultural heritage.

I would like to thank Wesley and Caro for proof-reading this document and supporting me. Additionally, I would like to thank Wesley for having patience with me and letting me talk about this project a lot in the last year.

I would also like to thank my family, especially my parents and my favourite brother Bogdan, for always supporting and encouraging me, not only during this project. Thank you as well to my friends who encouraged me and were curious and excited about my work.

Thank you to the science department team from the Rijksmuseum who allowed me to join their meetings and have a sneak peak into their amazing work and for their feedback on my prototype. Additionally, thank you to Francesca Gabrieli, Anna Krekeler and Nouchka De Keyser for letting me use the data on the two very beautiful paintings you are going to see in this report.

Lastly, thank you to all the people I met at the two conferences that I took part in during my thesis for their kind and inspiring words.

Andra Popa
Delft, December 2022

Contents

1 Introduction

Painting analysis reveals information about the materials used for the creation of a painting. This information can be used to understand an artist's work, for the conservation and restoration of paintings or for historical reasoning [1]. Generally, the painting analysis process starts with non-invasive techniques, in particular with imaging techniques that allow the identification of different pigments in the paint layers and the visualization of their distribution on the painting surface [2].

Two such non-invasive imaging techniques are Macro X-Ray Fluorescence (MA-XRF) and Reflectance Imaging Spectroscopy (RIS). MA-XRF reveals the chemical element distribution in paint layers and RIS reveals the molecular structure of the layers. However, analysing the results obtained using these two data modalities individually is not always enough for understanding the pigment distribution of a painting. MA-XRF and RIS provide unique insights into the structure of the pigments and complement each other [1]. Therefore, domain experts work with both data modalities to identify the pigments used in a painting.

MA-XRF and RIS are hyperspectral imaging techniques, which means that discrete values are registered for narrow bands and a continuous spectrum is obtained for every scanned pixel. The resulting data are saved as data cubes [3]. Such data cubes may contain millions of pixels with more than 200 dimensions per pixel. Domain experts are now facing the problem of efficiently exploring and interacting with these high-dimensional spaces. They identify pigments by analyzing the elemental and molecular profiles of pixels, which they compare to known pigment characteristics. A visual approach allows for easy exploration and interaction with the scanned pixels. However, directly visualizing and understanding the structure of the high-dimensional spaces of RIS and MA-XRF is challenging. To tackle this, we chose dimensionality reduction to visualize each dataset and make it interactively explorable. Furthermore, we use t-SNE[4] to project the datasets in two dimensions [5] and embed similar pixels as neighbouring points such that domain experts can relate groups of points to the pigments they represent. Additionally, the two data modalities are generally considered in separate steps of the analysis process [6, 7] and domain experts combine the gained insights from both data modalities without assistance from the tools they are using. Therefore, an integrated visual analysis system that connects the MA-XRF and RIS data can help domain experts better understand the elemental and molecular structure of pigments.

Furthermore, machine learning techniques exist for the automatic analysis of RIS data [8] that compare the scanned spectra to a library of existing pigment representative spectra to identify the pigments in a painting. However, it is difficult to model the complex mixtures of pigments in paint layers. Moreover, paints and painting techniques differ from one artist to another and from one time period to another, so there is limited comparability between different RIS datasets. Automatic techniques are limited to the spectra present in the spectral library, that may not represent an accurate baseline for the pigments in the painting. Domain experts can reason about the pigment distribution in a painting even in mixtures by looking for specific cues in one of the data modalities and connecting them to characteristics identified in the other data modality. Because of this, they are able to justify their decisions, unlike automatic methods, and modify their conclusions during the analysis as they uncover more information that is important. Therefore, domain expert knowledge is relevant even for automatic methods as experts can validate the accuracy of the automatic results.

We built an integrated system, VisuaLayered, which consists of multiple linked views for the exploration of both data modalities. VisuaLayered keeps domain experts in the loop by being highly interactive and allowing them to be involved in every step of the analysis process. We worked together with domain experts from the Rijksmuseum, Amsterdam in order to design, implement and evaluate VisuaLayered. The main contributions of this work are:

1. A combined and interactive visual analysis workflow for RIS and MA-XRF data with the goal of pigment identification
2. A prototype implementation of the mentioned workflow

We evaluated the efficiency of VisuaLayered in a case study with our collaborator. We wanted to see how well VisuaLayered implements the designed workflow, so we prepared five analysis tasks to be completed in VisuaLayered. Additionally, we asked our collaborator to complete the same tasks with existing tools that she normally uses in order to see if VisuaLayered improves her typical analysis workflow. The results show that the views combining the MA-XRF and RIS data provided in VisuaLayered represent a clear improvement in understanding the connections between the two data modalities. Furthermore, she was able to identify more pigments using VisuaLayered.

The project structure follows the design study methodology presented by Sedlmair et al. [9]. Design studies fall under the category of problem-driven research as they start with an existing problem and try to find a solution for it by first designing it and then validating it. Sedlmair et al. [9] propose a framework that leads to the implementation of a visual analysis system. The framework can be divided into three parts. The first part refers to the set up of the study in which we learned about the current visual analysis literature and we met with our collaborators from the cultural heritage domain. The second part consists of the abstraction of the problem, design, implementation and evaluation. The third and last part contains the reflection on the obtained results and the writing of the design study paper.

Chapter 2 presents the background of the two data modalities, MA-XRF and RIS and their application in cultural heritage. Chapter 3 highlights the related work with respect to current visual analysis systems for the two data modalities and for similar high-dimensional data. Then, in Chapter 4 we describe the requirements analysis and the process through which we defined it. In Chapter 5 we present the prototype, VisuaLayered, and explain how it connects to the requirements. Chapter 6 presents the evaluation and discussion of results. In Chapter 7 we present our ideas for future work and in Chapter 8 the conclusions.

2

Background

In this section we describe the two data modalities used by our collaborators, Reflectance Imaging Spectroscopy (RIS) and Macro X-Ray Fluorescence (MA-XRF) and explain how they can be used for painting analysis. Both data modalities are hyperspectral imaging techniques, which means that for every pixel of the scanned image, we register a continuous spectrum. This is possible because values are registered for neighbouring wavelengths with a small sampling distance [10]. Compared to multispectral techniques, which only register values for a small number of more distant wavelengths, hyperspectral techniques characterise a material in more detail and allow for easier comparison between materials [11]. Hyperspectral imaging analysis developed rapidly over the last 20 years and became available in the cultural heritage domain, where domain experts use it for pigment identification. Other applications of hyperspectral techniques include minerals' analysis in geology and soil and vegetation analysis [10].

2.1. Reflectance Imaging Spectorscopy

Reflectance Imaging Spectorscopy (RIS) determines the molecular composition of a material by measuring the amount of light that is reflected by it. Light interacts with elements and molecules [3] and we obtain a reflectance spectrum, where low values indicate that light is mostly absorbed and high values indicate that light is mostly reflected by the material. The devices used for RIS capture values in the visible range, 400 - 700 nanometers (nm) and in the infra-red range, 700 nm - 1 mm [12]. We can observe the entire electromagnetic spectrum in Figure 2.1, with values expressed in meters. Our collaborators work with VNIR (visual and near infra-red) cameras, which capture the range of 400 - 1100 nm and with SWIR (short wave infra-red) cameras, which capture the range of 1000 - 2500 nm.

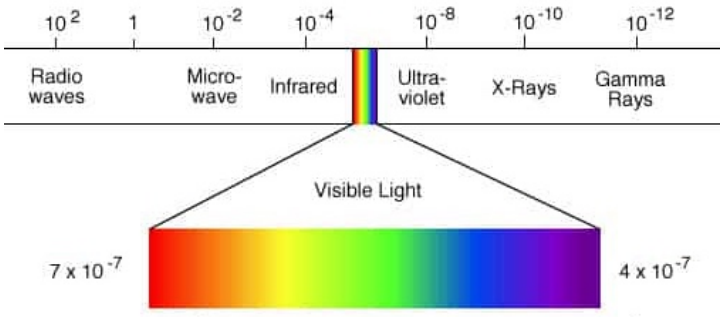

Figure 2.1: Electromagnetic spectrum (by William Anderson)

Every pure material can be identified based on its reflectance spectrum. We call this representative reflectance spectrum an endmember [10]. In an ideal world, this would allow us to identify any material by comparing its reflectance spectrum to a universal library of endmembers. However, in reality, materials (including the pigments used in paintings) are not pure due to their mixing and layering with other

materials or due to degradation because of aging or other external conditions. Furthermore, such a universal library of endmembers is often not available.

For painting analysis, domain experts may define endmembers as the representative spectra in the painting under study. Using spectral unmixing methods, we can identify the endmembers that constitute a mixed spectrum and their contribution to the mixture. There exist linear and nonlinear unmixing methods. Linear methods assume a spectrum is a linear combination of endmembers and provide accurate results when materials are present in the mixture in distinct layers. However, in cases where materials are mixed in layers that intertwine, nonlinear methods provide better results. Some linear spectral unmixing algorithms are based on three steps. The first one is dimensionality reduction in order to reduce the complexity of processing hyperspectral data, while preserving important information for the unmixing of spectra [13]. The second one is the identification of endmembers in the reduced space. The third one is the unmixing of all mixed spectra based on the identified endmembers. For the second step, we have two ways in which we can identify endmembers, manually or automatically. Automatic methods are used to identify endmembers directly in the scanned data or by matching scanned spectra to endmembers from a spectral library. Typically, the automatic methods that find endmembers in a painting, try to find an optimal set of endmembers according to a defined criterion, but such endmembers may no longer be physically meaningful. Therefore, it is important to still consider manual identification, which is done by domain experts, who use their domain knowledge of realistic endmembers and identify them by interactively exploring the scanned spectra and choosing representative spectra for different areas in the painting [13]. In some cases, endmembers have peaks at specific wavelengths, which clearly indicate the presence of a particular pigment. Alternatively to using unmixing methods, we can use a spectral similarity measure in order to identify the spectra that are similar to the identified endmembers. By considering only the spectra that are similar to an endmember, we obtain an endmember mapping, which shows us the spatial distribution of relevant molecules.

2.2. Macro X-Ray Fluorescence

X-Ray spectroscopy appeared at the beginning of the 20th century. Macro X-Ray Fluorescence (MA-XRF) determines the elemental composition of a material. This technique works by shooting an X-Ray that moves an inner shell electron of an element and then registering the fluorescent radiation created when this electron is replaced by an outer shell electron [15]. Figure 2.1 shows that X-Rays are situated between ultraviolet and Gamma rays on the electromagnetic spectrum. The peaks of the XRF spectrum at specific energy values indicate the presence of different chemical elements (Figure 2.2). Therefore, based on these peaks, experts extract elemental maps showing the distribution of one chemical element and further study the elemental composition of a material.

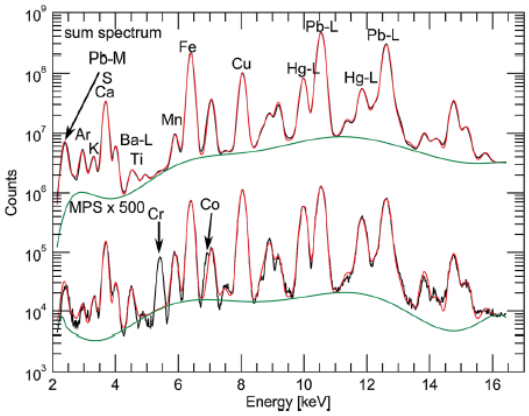

Fig. 7 Fit of sum spectrum and Maximum Pixel Spectrum. In the fit model Cr, Co, As and Bi are not included. Spectra fitted with PyMCA.

Figure 2.2: Example XRF spectra [14]

2.3. RIS and MA-XRF Data in Cultural Heritage

Paintings have a complex, layered structure, which we can understand through advanced imaging techniques like RIS and MA-XRF. A painting is created on a support (e.g. a canvas) and consists of a ground layer, followed by an underdrawing, pictorial layers and lastly, a varnish layer [16]. The pictorial layers contain different pigments, that domain experts try to identify. If they know the pigments and the order and quantity in which they are applied, domain experts can accurately restore and conserve paintings, without altering their composition. Through non-invasive scanning techniques, like RIS and MA-XRF, domain experts may even discover a completely different painting underneath the visible paint layers [17] or observe the corrections done in the past by other restorers. Moreover, they know how pigments alter over time and can identify specific cues in the data that indicate the age of the used

pigments. Additionally, art historians gain insights into the pigments used in different time periods by important artists [18] through the analysis of multiple, varied paintings.

Non-invasive imaging techniques, such as MA-XRF and RIS, are the first part of a typical analysis process. Galli and Bonizzoni [19] make a summary of various XRF techniques and their applicability in cultural heritage. They present MA-XRF as a powerful and easy to use technique that can guide a more in depth analysis of a painting. If necessary, domain experts take samples of a painting and do experiments on it in order to have a complete analysis of a painting. Mazzinghi et al. [20] used MA-XRF to analyse a painting of a Flemish painter from the 15th century. This technique alone could not identify pigments with 100% certainty, but in combination with results from other works from the same painter or from the same time period, they could still decide on a set of possibly used pigments. With the use of MA-XRF element mappings they identified pigments like ultramarine, vermilion and gold. Moreover, they discuss the use of MA-XRF for the analysis of furniture, tiles, metals, manuscripts and other materials in cultural heritage, highlighting its relevance in the field. Applications of RIS also reveal important information about the paint layers in a painting [21, 22]. Recent works focus on automatic endmember detection. For example, Grabowski et al. [8] successfully identify the endmembers of five green pigments by automatically matching the scanned spectra to 22 endmembers in a spectral library. They propose a three-step method which consists of clustering the data, identifying endmembers based on the clustering followed by mapping the endmembers to the spectral library and in the end, labeling the pixels based on the pigments they represent. However, a limitation of their results is the fact that they only test their method on a small number of pigments in mock-up samples and not in an artist's painting, which may have aged pigments in more complex mixtures, that may not be represented that easily in a spectral library. Moreover, there is work focused on finding accurate spectral similarity measures for automatic matching [23]. Additionally, Grillini et al. compare several imaging models for spectral unmixing [24]. Their models fall into three categories: additive, subtractive and hybrid. They conclude that subtractive mixing models perform the best in the case of oil-based mock-up paintings.

MA-XRF and RIS are used in combination in the cultural heritage field as they provide complementary information that is useful to better understand the structure of a painting. They allows us to uniquely identify pigments and visualize their distribution on the surface of a painting. On the one hand, MA-XRF indicates the presence of different chemical elements (e.g. calcium, copper) in the paint layers of a painting, but this is not enough information for pigment identification. Firstly, one chemical element can be present in multiple pigments. Secondly, paint is applied in layers and MA-XRF does not provide depth discrimination, so we may not always know how the elements are distributed in the layers. Thirdly, pigments can be mixed, so it may be unclear which elements belong to which pigment. On the other hand, the RIS data also presents limitations. Some pigments have very similar endmembers, so it is difficult to distinguish between them only based on their measured reflectance. Moreover, if we have a mixture of pigments, its endmember may not be enough for identifying its constituent pigments. Therefore, we can use the RIS data in combination with the MA-XRF data in order to characterize pigments with respect to both their molecular and elemental compositions in order to make an informed decision on the pigment distribution in a painting.

Additionally, the two data modalities have different capabilities with respect to the interaction with the paint layers, as they have different depths of penetration that depend on the energy of the used radiation. MA-XRF provides information about all the layers of a painting because it penetrates the entire thickness of an object, while RIS only interacts with the upper layers of paint because light only penetrates the first microns. Hence, a scanned pixel might indicate the presence of chemical elements that do not correspond with the RIS results. Hence, domain experts have to reason about the possible combinations of chemical elements that match the identified endmembers. Moreover, they can deduce what elements are present in the first layers of paint, that are not covered by RIS and have a clearer idea of the order in which the pigments were applied in the painting.

As recent works show, the combined analysis of MA-XRF and RIS proves to be very powerful in the cultural heritage domain. Domain experts work with MA-XRF elemental maps (Figure 4.1b) and RIS endmember maps (Figure 4.1c) in order to reason about the pigment distribution in a painting. For example, D'Elia et al. [25] conducted an analysis on two paintings made by one of Leonardo da Vinci's apprentices. They scanned one section for each painting as they were interested in modifications made in those areas. By using both data modalities they successfully identified pigments like carbon black and vermilion. Moreover, they identified parts of the original painting, that were hidden underneath the top paint layers. In another example, the combination revealed pigment residues on an antique Greek

statue [26]. Even though the statue was extensively analysed in the past by archaeologists, the two data modalities complemented each other and showed the presence of earth based underlayers of paint and revealed the presence of lead based pigments on the statue. Additionally, the combined analysis revealed in high detail the pigment distribution in Johannes Vermeer's *Girl with a pearl earring* [27]. For this application, MA-XRF was used to confirm the pigments found based on the endmembers (e.g. for the pigment vermilion, MA-XRF showed the presence of mercury) and to reveal information about the sub-surface layers that were not captured with RIS. Moreover, the combination revealed the distribution of the pigments identified in samples taken from of the painting in previous analyses.

3

Related Work

We are interested in the combined visual analysis of RIS and MA-XRF data. Since these two data modalities are hyperspectral, we start by looking into existing analysis techniques for hyperspectral data in Section 3.1 in order to learn how they are typically analysed and develop our workflow based on them. Moreover, due to the fact that domain experts try to establish the pigment distribution in a painting by exploring different areas of the painting, we consider visual analysis systems that allow users to interactively explore the search space in section 3.2. Additionally, due to the high-dimensional nature of the two data modalities, we consider dimensionality reduction techniques that project the data into a structured lower-dimensional space that domain experts can more easily interact with.

3.1. Analysis of Hyperspectral Data

The existing software for the analysis of hyperspectral data for cultural heritage generally focus on the analysis of one of the two data modalities used in this work. XISMuS is a software system designed specifically for MA-XRF data analysis [6]. It provides pre-processing methods for MA-XRF data cubes. Moreover, it supports the visualisation of elemental maps, individually and in combination, similarly to our system. In addition, we provide the visualization of combined maps from RIS and MA-XRF such that users can explore the spatial overlap of the two data modalities. PyMca is another software system for the analysis of XRF data [28]. Campos et al. [29] use PyMca for creating and visualizing elemental maps. While at the beginning the system focused on one type of XRF maps, it developed over time and now also supports the imaging of other (even non-spatial) multispectral data, like Fourier transform infrared maps. Moreover, it supports the simultaneous manipulation of different data modalities [30]. However, PyMca does not provide a dedicated workflow for pigment identification, which is our goal.

Next, we consider systems that allow users to explore the reflectance spectra in RIS datasets. The Specim IQ hyperspectral VNIR camera comes with its own analysis software [7]. For a scanned painting, users can select individual pixels in order to visualize their spectrum. Moreover, users can create false color images and explore the spatial distribution of the selected spectra using the Spectral Angle Mapper (Section 5.3) algorithm. We provide similar functionalities, with improved interactivity, such that we can define and save endmembers based on selections of multiple similar pixels. Additionally, HypIX (Hyperspectral Image eXplorer) allows us to explore time series of hyperspectral data [31]. Through three interactive, linked views they show the spatial distribution of data and the registered spectra at different time steps and an embedding of all time steps. Our system consists of similar linked views, but instead of focusing on how the hyperspectral data evolves over time, we focus on visually combining the two data modalities in order to better understand the correlation between them. Furthermore, Hyper3D is a system that focuses on the visualization of volumetric hyperspectral data [32]. It was designed in collaboration with art conservators and is meant to help them in their analysis. They provide 2D views to complement 3D models' inspection, like a spectral plot showing the spectrum of a selected pixel. We follow a similar design process, but work with linked 2D visualizations of hyperspectral data and provide a spectrum view with more interactive settings for endmember identification.

Furthermore, since hyperspectral data is relevant for other fields, such as agriculture and mineralogy [24], tools developed for these fields can still be useful for cultural heritage data analysis. Such

an example is the geospatial data analysis software, ENVI [33]. It is a powerful system that provides many functionalities that can be used for analysing RIS data, including the ones present in the aforementioned Specim IQ software. Additionally, ENVI allows users to automatically detect endmembers using the Spectral Hourglass Wizard. We describe ENVI in more detail in Section 4.3, as it is the system that our collaborators currently use for their analysis.

For our implementation, we integrate views and techniques that are used in the field for each data modality, such as the ones described in this section. However, the presented systems are not designed for multi-modal data analysis. We want to go a step further by designing and implementing a workflow that focuses on the combined analysis of the two data modalities.

3.2. Interactive Visual Analysis of High-Dimensional Data

The data obtained using hyperspectral imaging can be explored through visualization. Liu et al. [34] conducted a survey into high-dimensional data visualizations. We can directly visualize the data using Parallel Coordinate Plots (PCP), which treat the dimension of a dataset as separate axes, ordered linearly and represent one point as a line that has a specific value for every axis. This view allows us to see points that have similar/different values across dimensions, but the structure that we observe is influenced by the ordering of the axes. Another direct visualization method is the Scatterplot Matrix (SPLOM) which consists of a matrix of scatterplots showing the relation between every two possible dimensions of a dataset. However, this view does not allow for exploring the relationship of more than two dimensions at the same time.

The high-dimensional nature of the data makes it difficult to directly visualize it. We can address this by transforming the original data through subspace clustering, dimensionality reduction (DR) or regression analysis [34]. Dimensionality reduction (DR) projects high-dimensional data into a lower-dimensional space, typically 2D or 3D, while preserving important features between points in the original space. We chose dimensionality reduction in our work because we wanted a space that that can be efficiently explored and interacted with. DR techniques can be divided into two main categories, linear and non-linear. Linear methods, such as Principal Component Analysis (PCA) [35] and Linear Discriminant Analysis (LDA) [36] transform the data linearly. Non-linear methods, such as Uniform Manifold Approximation and Projection (UMAP) [37] and Multidimensional Scaling (MDS) [38] create projections for data that cannot be modelled linearly [34]. Another non-linear method is t-SNE [4], which preserves local distances between data points when performing the projection. We use t-SNE in our analysis, as domain experts are interested in identifying pigments, which correspond to points that are similar in the high-dimensional space with respect to both data modalities. t-SNE has many applications in medical visualization, for example in the case of singe-cell data analysis, it allows the identification of potentially problematic tissues [39]. Moreover, in medical genetics t-SNE can separate continental populations and accurately represent human genetic data at different scales (global and local) [40]. In the case of the analysis of high-dimensional images, Vieth et al. [41] incorporated texture information into dimensionality reduction methods.

Dimensionality reduction is already used for the analysis of hyperspectral data [11]. In the case of ink and paper data analysis, the t-SNE embedding shows a clear separation between groups of similar points and facilitates the easy clustering of the data, outperforming PCA due to its non-linearity [42, 43]. The embeddings are visualized in 2D scatterplots. Moreover, the clusters obtained based on the t-SNE embedding score better than the PCA based ones according to different clustering indices. Pouyet et al. [44] conduct a comparison between the same two algorithms and minimum noise fraction (MNF) for the reduction of reflectance data of mock-up paintings and one manuscript. In this case, t-SNE also produces a better embedding for endmember identification through clustering. In HypIX [31], which is mentioned in Section 3.2, the embedding view obtained using t-SNE acts as a basis for point selections and further analysis in the other two available views. Additionally, in Section 2.3 we mentioned an example of automatic endmember identification [8]. As part of their algorithm, they use dimensionality reduction (Barnes-Hut-SNE [45]) for creating an embedding fit for the automatic clustering of pixels. Macro et al. [25] use PCA for filtering pixels that are perceptually similar, but different in their reflectance spectrum, when trying to find representative pixel clusters for endmember identification. Additionally, using the t-SNE embedding of a joint data cube of MA-XRF and RIS data, Alfeld et al. [3] could more easily establish some of the pigments present in the paintings of an Egyptian tomb. For example, the embedding showed a clearly separated cluster for one pigment and highlighted sub-surface paint

layers. The combination of the two data cubes is non-trivial, as they span different high-dimensional spaces and have to be normalized and perfectly aligned in order to obtain a fused cube that can be used for meaningful further analysis. Based on previous work, we see that t-SNE embeddings facilitate the easy and accurate clustering of hyperspectral data.

In the astronomy field, Goodman highlights the need of a complete system for the analysis of high dimensional data [46]. He suggests the use of multiple linked views that update upon a user selection, not only for tabular data, but also for images and data cubes. Such a system would enhance the analysis process by offering context in a visualization, which allows for comparison between different data models. He also mentions integrating views available in the field and adapting views from other fields in order to build such a system. We aim to create a similar integrated system for the combined analysis of RIS and MA-XRF data.

Interactivity is an important element in the analysis of high-dimensional data as it allows domain experts to focus on specific aspects of the data and to integrate their knowledge into the system in order to solve their analysis tasks. Brehmer et al. [47] focus on the visualization of dimensionally-reduced data and make a classification of visual analysis tasks based on interviews with ten domain experts from different fields. Three of the five identified tasks refer to the clustering of the reduced space and are suited to our workflow. The first task is identifying clusters in the reduced space and learning more about the structure of the high-dimensional space based on those clusters. The second task is naming the clusters in order to give them meaning and to summarize the analyst's results. In VisuaLayered, pixels that are similar with respect to either data modality may be clustered and named according to the pigment they represent. The third task refers to the matching of clusters to the classes they represent. Moreover, Sedlmair et al. [48] analyse DR uses based on the problems experts from different domains try to solve. In many data analysis cases, DR is used for creating an embedding that can be visualized and explored in a 2D scatterplot. They divide the user tasks into dimension and cluster focused tasks. Sacha et al. [49] mention the importance of adapting DR based analyses to domain specific problems by involving experts in the design process and providing interactive visualizations that fit their analyses.

The analysis of high-dimensional data from different domains is supported in interactive visual analysis systems. Cytosplore is a visual analysis system for the analysis of the high-dimensional mass cytometry data, which allows the user to interactively explore the data in order to identify and annotate phenotypically similar cells [50]. ImaCytE extends the notions presented in Cytosplore for the analysis of imaging mass cytometry data [51]. One of the tasks supported by ImaCytE is the identification of cell phenotypes. The user can visualize the data in a 2D embedding, cluster similar cells based on the embedding and annotate the found clusters. The clusters can be iteratively changed by the user and changes have immediate effects in all the other linked views of the system. We follow a similar workflow for identifying endmembers, but in addition, we combine and correlate the endmember clustering with the MA-XRF clustering in order to relate the molecular structure of pixels to their elemental structure and identify pigments. Another integrated system for the analysis of imaging mass cytometry data is histoCAT [52]. Furthermore, using Visinity [53] users can visually analyse cell interactions by considering cell neighbourhoods in tissue images. The user can select regions of interest in an image viewer, further explore local neighbourhoods in a parallel coordinate plot, explore similar neighbourhoods in a UMAP 2D embedding and explore cell interactions in a correlation matrix. Similarly, we combine MA-XRF and RIS data in multiple linked views. Visinity builds on two previous systems, Facetto [54], which supports the partially automated analysis of tissue cells and Scope2Screen [55], which facilitates focus+context exploration and annotation for the analysis of tissue images. Additionally, the visual analysis workflow proposed by Somarakis et al. for the analysis of spatially-resolved omics-data allows users to identify differentiating characteristics of two cohorts in order to better understand the evolution of diseases [56]. They use customizable raincloud plots in order to visually compare the two cohorts based on the abundances of different cell types. They focus on identifying the differences between two datasets, while we try to identify the similarities between two data modalities since pixels that represent one pigment or a mixture of pigments should have the same composition with respect to both data modalities. Visual analysis is also beneficial for exploring temporal data. MulteeSum is a system that visually combines temporal gene expression data with spatial data of the cells that contain the explored genes [57]. Lastly, EVA is a visual system used for identifying financial fraudulent operations based on money transaction data over a period of time [58].

4

Requirements Analysis

In order to develop a visual analysis system for painting analysis, we needed to understand what the requirements of such a system are. As described in Chapter 1, we structured our work as a design study and in this chapter we describe all the steps that were required to start the implementation of our prototype. In Section 4.1 we give an overview of the entire design process from learning about the use of MA-XRF and RIS data for painting analysis to evaluating VisuaLayered. In Section 4.2 we describe the data specifications for MA-XRF and RIS using a painting provided by the Rijksmuseum as an example. Then, we present the typical analysis workflow of our collaborator in Section 4.3 and based on it, we describe the task abstraction for our proposed workflow in Section 4.4 and the visualization design for VisuaLayered in Section 4.5.

4.1. Design Process

This design study was carried out in close collaboration with imaging researchers at the Rijksmuseum, Amsterdam. We started the study with two introductory meetings with our collaborators. In the first meeting, we were introduced to the two data modalities that we use in this project, MA-XRF and RIS. Moreover, we briefly observed the current workflow of our collaborators. In the second meeting, they presented us general acquisition methods and their applications in the cultural heritage domain. Additionally, we saw in more detail how MA-XRF and RIS work and what kind of information they capture about a material. These meetings helped us understand the background for our work.

Afterwards, we conducted a field study in order to learn more about the typical analysis workflow of our collaborator, so that we understand how we can help them with a combined visual analysis system. For the field study, we had the opportunity to visit the Rijksmuseum and observe our collaborator while they analysed the painting described in Section 4.2 using MA-XRF and RIS data. Based on the field study, we created a visualization design for VisuaLayered, which we presented to our collaborator in order to make sure we included all of their requirements. Next, we deployed the first prototype of VisuaLayered. Moreover, we sent our collaborator different versions of the prototype once new features were added such that they can test it along the way. Throughout the project, we had regular meetings with them during which we updated them on our progress and we received feedback on the different functionalities of VisuaLayered.

4.2. Data Structure

The main painting that we used during the development of the prototype is called *Figures in a Courtyard behind a house* by Pieter De Hooch [59], which can be seen in Figure 4.1a. Pieter de Hoch was a Dutch Golden Age painter, who created this work in c.1663-c.1665 in his Delft period. The painting is oil on canvas and has a size of 45.7 x 60 cm (width x height). Data on the painting was provided by our collaborators from the Rijksmuseum. The MA-XRF data was acquired and processed by Anna Krekler [59] and the RIS data was acquired and stitched by Francesca Gabrieli. For MA-XRF, the data was captured with a Bruker M6 scanner. We received elemental maps, such as the one in Figure 4.1b, which shows the distribution of Ca in the painting. For RIS, our collaborators shared a dataset obtained using a SOC710 VNIR camera. Based on the RIS dataset, we can obtain endmember maps, such as

the one shown in Figure 4.1c. For the RIS data, the painting was scanned in horizontal bands using a spectral sampling distance of 2.54 nm (in a similar setting as shown in Figure 4.2). In order to obtain a full image of the painting, the bands were stitched back together.

Figure 4.1: a) *Figures in a Courtyard behind a house* by Pieter De Hooch [59], b) *Ca* elemental map, c) *Red ochre* endmember map

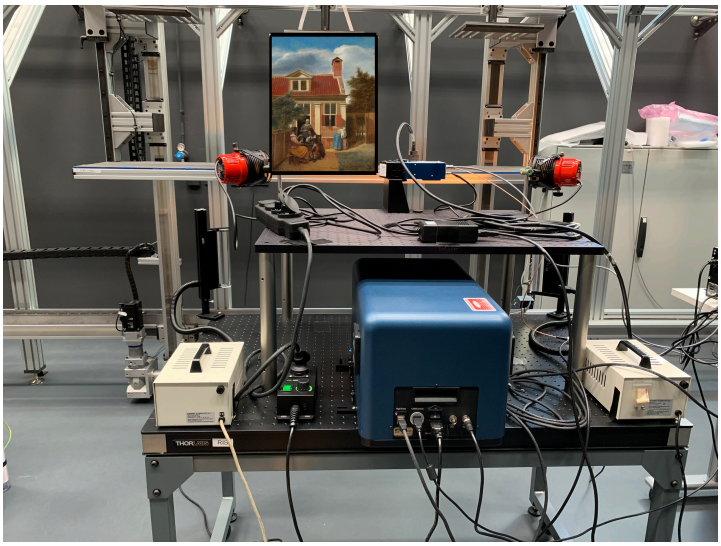

Figure 4.2: VNIR camera example setup

As mentioned in Chapter 2, RIS and MA-XRF are hyperspectral imaging techniques, so for both data modalities we register a data cube which contains one spectrum per pixel. For RIS, we consider a discretely sampled version of each reflectance spectrum that we interpret as a high-dimensional space. RIS can have a spectral resolution of less than 10 nanometers and a spatial sampling of hundreds of micrometers [60]. The raw MA-XRF data cube may contain millions of pixels with thousands of dimensions per pixel [14]. With mobile and commercial MA-XRF devices, it is possible to scan surfaces of $\approx 50 \times 50$ cm^2 in a few hours, with lateral resolutions ranging from tens to thousands of micrometers. However, after the raw cube is processed into elemental maps, the dimensionality is reduced to the number of identified chemical elements. One elemental map represents one dimension of the final MA-XRF dataset. The number of obtained maps is ≈ 10 [61]. Hence, we work with RIS datasets that have a higher number of dimensions than the MA-XRF datasets. For both datasets, the x and y coordinates indicate a pixel's position in the painting and the z coordinate indicates its dimension. Figure 4.3 and

Figure 4.4 show illustrations of a RIS and an MA-XRF data cube with spectra for two extracted pixels, respectively.

As a requirement in our analysis, the two datasets must have the same number of high-dimensional pixels, such that we can characterize a pixel with respect to both data modalities. We realise that for the best results the two datasets have to be completely aligned, such that the pixels represent the same location in the painting for both data modalities. However, such data was not available during the development of the project and we discuss this limitation further in Section 6.3.

The painting shown in Figure 4.1a is saved as 1174 × 1756 pixels for both data modalities with a pixel size of roughly 0.35 × 0.35 mm^2 (after re-scaling). For RIS, we received a sampled data cube, which has 200 dimensions in the range $[400, 900]$ nm. For MA-XRF, we received 28 grayscale element mappings, each showing the spatial distribution of one chemical element in the painting.

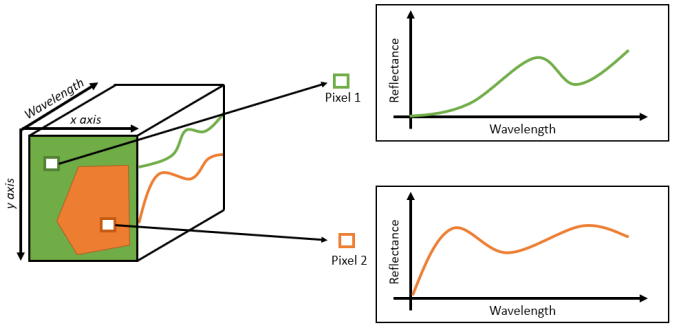

Figure 4.3: RIS data cube

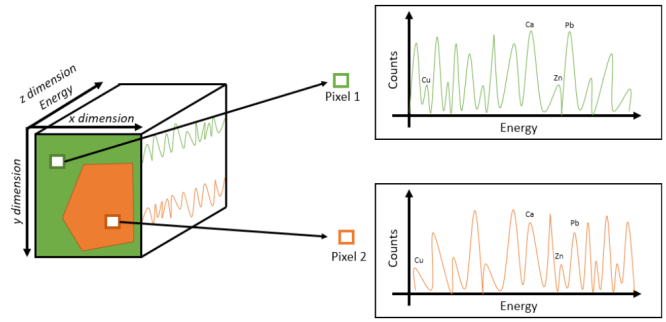

Figure 4.4: MA-XRF data cube

4.3. Field Study

We set up the field study with our collaborator in order to learn about her typical analysis process. It was important to first understand the steps of the workflow for the analysis of MA-XRF and RIS data such that we can design our combined workflow. We mostly observed her workflow and asked questions when it was unclear to us why she took certain decisions in the analysis process. Her analysis normally starts with the visual inspection of the MA-XRF maps to get an initial overview of the elemental distribution in the painting.

Then, she uses the geospatial analysis software ENVI [33] for analysing the RIS data. She uses three main views provided by ENVI for identifying endmembers in the painting: an image viewer, a lineplot and a scatterplot. Using the image viewer she can visualize an image of the painting, either in grayscale by choosing one wavelength of the data cube or in RGB, by choosing three wavelengths for the red, green and blue channels of the image. The image view comes with 2 zoom levels for more precise pixel selections. She initially explores the spectra in the image by selecting pixels in different

regions in the image viewer and visualizing their spectra in a lineplot. In order to get a comprehensive set of endmembers, our collaborator identifies endmembers manually and automatically. For the manual selection, she searches for endmembers in a reduced space of the data cube. The Spectral Hourglass Wizard (SHW) algorithm, which is a variant of PCA, projects the data in a 3D scatterplot based on three chosen principal components. In this space, our collaborator manually selects pixels and looks at their spectra in the lineplot. If she finds a group of points with similar spectra that all represent one endmember, she saves the group as a cluster with its corresponding average spectrum as the endmember. After this, she also computes endmembers automatically and manually verifies whether the identified endmembers have a realistic spectrum.

In the end, for each endmember, she maps the pixels with spectra similar to it. The mapping is done using the Spectral Angle Mapper (Section 5.3). This algorithm uses a threshold value for defining the similarity between points, that our collaborator interactively chooses based on a histogram. In order to decide the final threshold value, she selects pixels in the image view that are included in the map and look at their spectra. For example, if some of the spectra are not similar enough to the endmember, she reduces the threshold value. Moreover, she already considers the presence of a pigment based on the endmember map and check it by reopening certain elemental maps that should also indicate the presence of the same pigment. Combined with the MA-XRF elemental maps, she has mappings for both MA-XRF and RIS and can visualize them side-by-side in order to identify the pigments used in the painting.

For the RIS data, in the lineplot, she can also interactively choose three wavelengths for the RGB channels of the image in order to obtain false colour images. Such images can help distinguish pixels which are spectrally different, but perceptually the same by for example, mapping one of the RGB channels in the near infra-red range. Moreover, some pigments are more difficult to distinguish by just using the measured spectra, so she performs the same analysis process on a first derivative data cube. This is especially useful in the case of red pigments. After this meeting, we considered the design of our system.

4.4. Task Abstraction

Based on the field study, we define a set of analysis tasks for the proposed system. In order to describe the tasks, we use keywords presented by Brehmer and Munzner [62] for the task abstraction of a visualization design, which are highlighted in *italics*. The two main tasks are the endmember identification based on the RIS data and the pigment identification based on the combination of MA-XRF and RIS data. We further divided these two tasks into subtasks that focus on specific actions (the keywords) that are required in order to complete the tasks:

T1. Endmember identification in RIS

- **T1.1** *Identify* candidate endmembers in the high-dimensional space
- **T1.2** *Compare* and *filter* candidate endmembers based on their spectral characteristics
- **T1.3** *Explore* the spatial distribution of pixels that are similar to a selected endmember

T2. Pigment identification

- **T2.1** *Explore* the spatial distribution of pixels containing selected chemical elements
- **T2.2** *Identify* pixels that have similar molecular and elemental profiles
- **T2.3** For the pixels spectrally similar to an endmember, *identify* their elemental composition
- **T2.4** *Annotate* pixels according to the pigment they represent, based on both data modalities

The tasks can be followed in the order in which they are written down (Figure 4.5) and hence the pigment identification can be seen as the final goal. However, it is important to note that the second task, the pigment identification, is strongly connected to the first one. Hence, some tasks may be completed at the same time in the workflow. For example, during task **T1.3**, domain experts may also consider **T2.1** in order to verify their assumptions regarding the identified endmembers. Moreover, since during the analysis, users might reconsider their endmembers and pigment classifications, our workflow allows users to easily go back to completed tasks and change their selections. We highlight three such

cases in Figure 4.5. Domain experts might reconsider their endmembers once they look at the spatial distribution of similar pixels (**T1.3** to **T1.1**) or once they identify pixels that are similar with respect to both data modalities (**T2.2** to **T1.1**). Additionally, when establishing the elemental composition of an endmember they might go back to analysing the elemental maps in more detail (**T2.3** to **T2.1**).

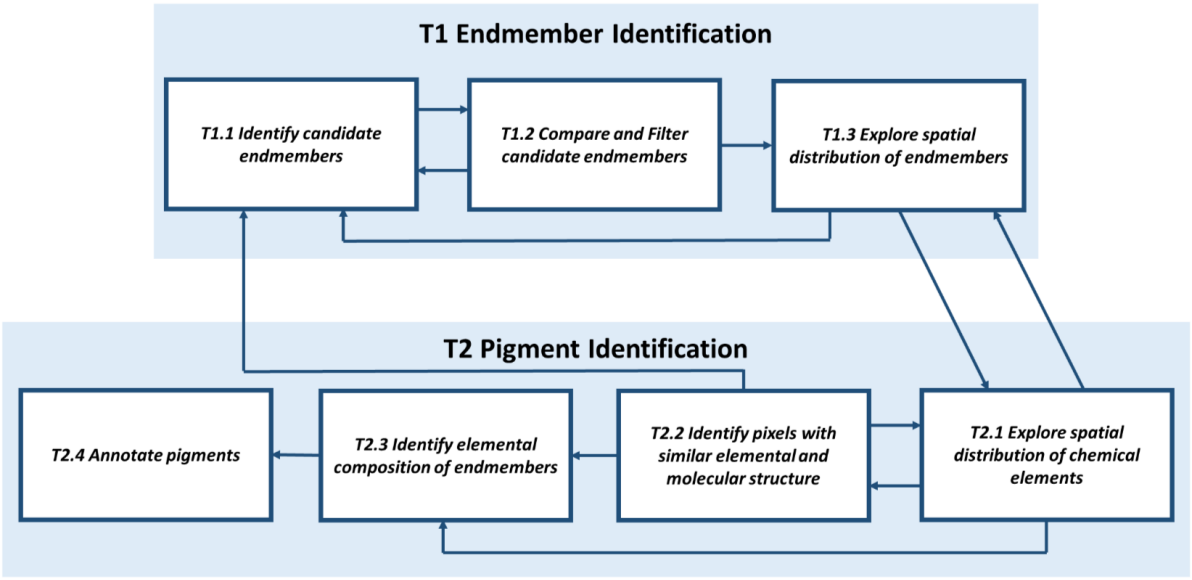

Figure 4.5: Workflow

4.5. Visualization Design

After learning about the workflow of our collaborators during the field study and establishing the required analysis tasks, we created a visualization design for our system. Our focus was on the combined visual analysis of the two data modalities, so we thought of different ways in which their connection can be explored. Our collaborator's current analysis is quite time consuming as there is no visual link between the two modalities and they combine the datasets mostly in their head. Moreover, we also wanted to incorporate parts of their current workflow regarding the analysis of the RIS data and improve certain interactivity aspects. In comparison to ENVI, which was not developed for cultural heritage analysis, we designed a system that is streamlined to painting analysis. Figures 4.6,4.7 and 4.8 show the initial design sketches, that we presented to our collaborator. We based our design on three possible use cases, that present all the initially considered views and ways in which they can be combined during the analysis. In total, we present three views: an image view, an embedding view and a spectrum view. Selections can be linked across all views so that we are always aware how pixels are represented in both data modalities. The three uses cases are:

1. **Combined cluster analysis** (Figure 4.6) - The first use case refers to the combined analysis of the two data modalities. The first view on the left is an image view, similar to the one provided in ENVI. The next two views are embedding views, which show the reduced space of the RIS and MA-XRF data. One point in the embedding represents one pixel. Compared to the Spectral Hourglass Wizard, our embedding computation requires less parameter settings. We provide the embedding view as a base for clustering points that are similar. In the case of the RIS data, it can be used for identifying endmembers and in the case of the MA-XRF data, it can be used for identifying pixels that have the same elemental profile. An example selection is shown in green, synchronized in all views.

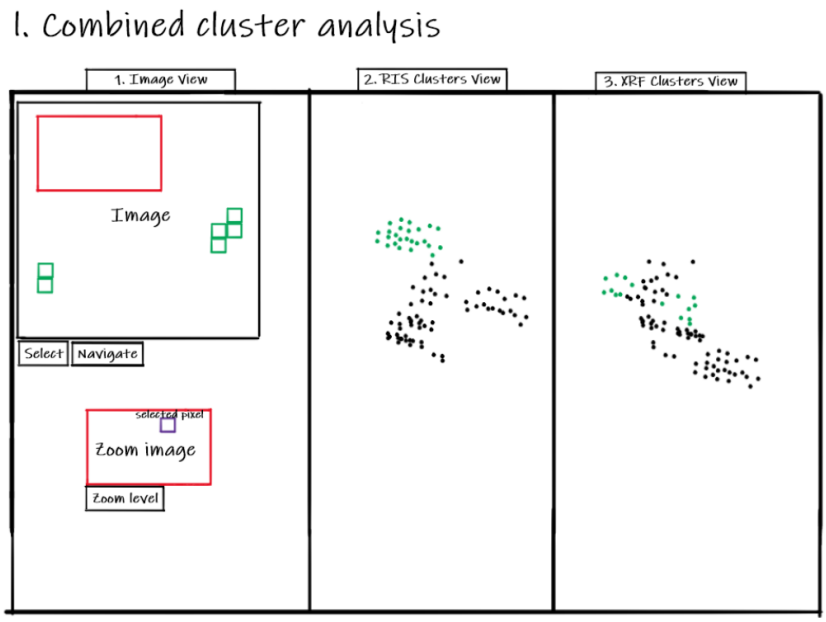

Figure 4.6: Combined cluster analysis design sketch

2. **Endmember selection** (Figure 4.7) - The second use case refers to the identification of endmembers based on the RIS data. The first two views are the image and embedding views. The third one is the spectrum view, that was designed based on the lineplot view in ENVI. We can see the three RGB lines in the spectrum view that can be connected to the RGB channels of the image. It is also possible to create endmember maps in this view. We also thought of ways to improve the interaction with the spectrum view. We show the standard deviation for the average spectrum of a selection of pixels. This addition informs the user how accurate an endmember is for a group of pixels. Additionally, dragging the RGB lines in the spectrum view and changing the threshold value for the endmember map computation, show results in real time such that the user receives immediate feedback on their actions.

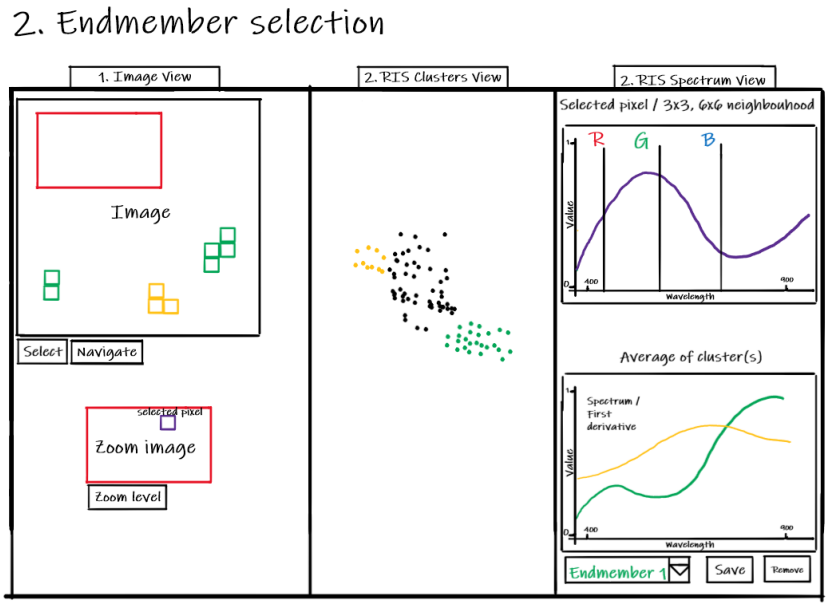

Figure 4.7: Endmember selection design sketch

3. **Mapping analysis** (Figure 4.8) - The third use case refers to the analysis of the two data modality mappings. Using the image view with three different inputs (an image of the painting, RIS maps and MA-XRF maps) we can reason about the spatial overlap of pixels with selected molecular/elemental structures. The maps can be visualized individually or in combination.

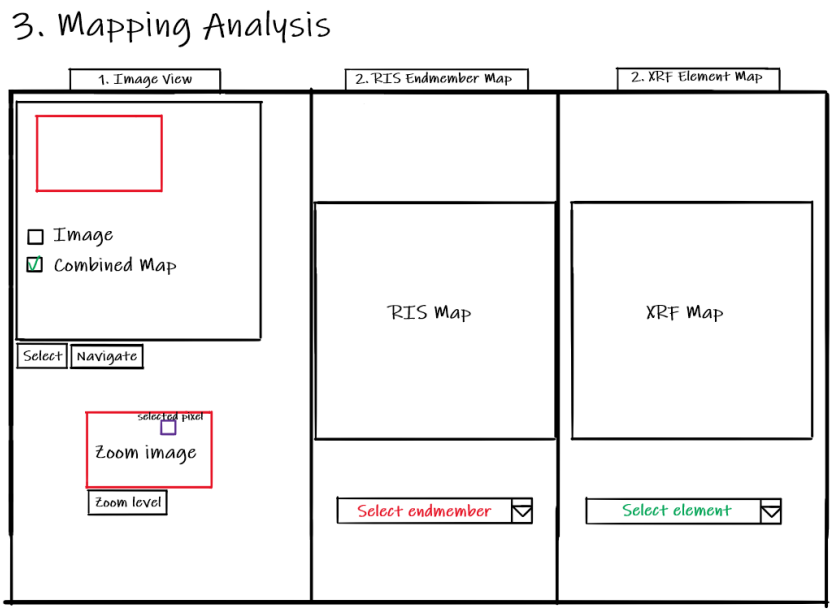

Figure 4.8: Mapping analysis design sketch

At this point in the design process, we also considered a view that shows the correlation between the clusters of the two datasets. This view is not included in the sketches, but they were added to the system during development.

5

VisuaLayered

VisuaLayered is the prototype that we designed and implemented for the combined visual analysis of MA-XRF and RIS data. The main building blocks of VisuaLayered are five interactive, linked views. Figure 5.1 shows a screenshot of VisuaLayered. The user can open the available views in any number, combination and order, to help them in their analysis. We built the prototype based on the tasks described in Section 4.4. In VisuaLayered, we can explore the spatial distribution of pixels for **T1.1**, **T1.3**, **T2.1** (Figure 5.1a), classify similar pixels for **T1.1**, **T2.2** (Figure 5.1b,c), compare pixels' reflectance spectra for **T1.2** (Figure 5.1d), explore the correlation between the two data modalities for **T2.2**, **T2.3** (Figure 5.1e) and annotate groups of points based on the pigments they represent for **T2.4**. Moreover, we can connect the two data modalities and explore pixels with respect to their molecular and elemental composition in order to identify the pigments used in a painting. Selected pixels in VisuaLayered are highlighted in red in all the shown images.

Sections 5.1 and 5.2 present the prototype functionalities and analysis workflow in relation to the tasks in Section 4.4. Section 5.1 presents the workflow for the endmember identification task (**T1**), using the RIS data. This subsection is based on the contents our short paper attached in Appendix A. Section 5.2 presents the combined analysis of the two data modalities for the pigment identification task (**T2**). In the last section of this chapter, Section 5.3, we describe the implementation of VisuaLayered.

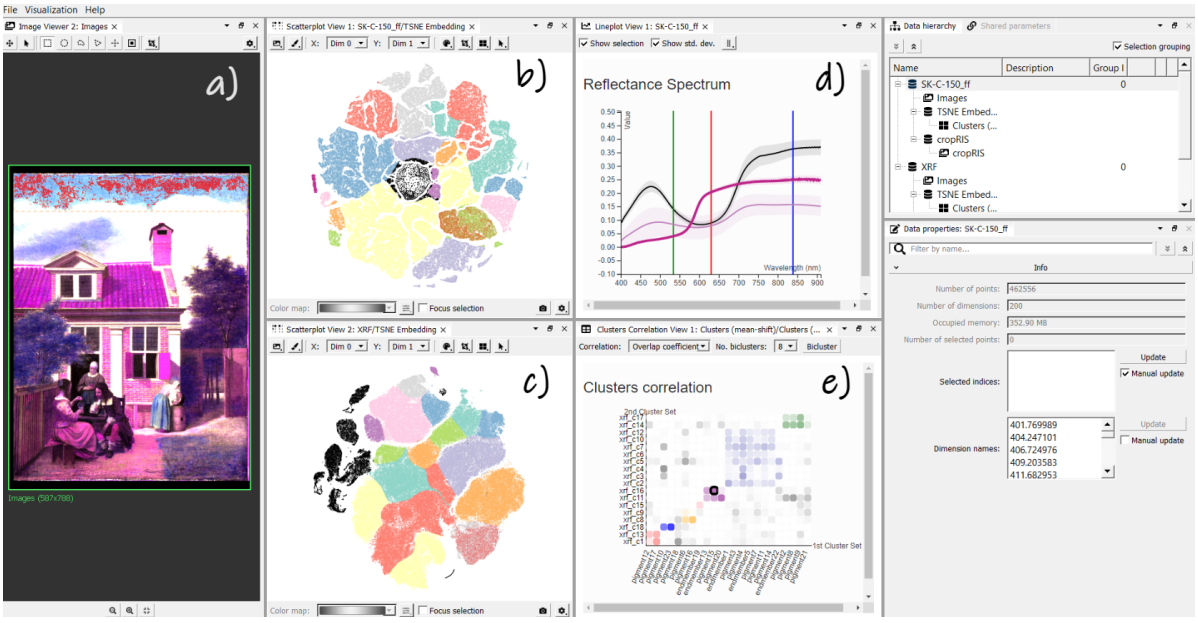

Figure 5.1: VisuaLayered showing the painting *Figures in a Courtyard behind a House* by Pieter De Hooch [59]; a) Image View showing an RGB image of the painting using the RIS dataset; b), c) Embedding views showing both data modalities; d) Spectrum view showing spectra from the RIS data; e) Cluster correlation view showing the correlation between the two data modalities

5.1. Endmember Identification

Here we present our design and implementation to tackle task **T1**, the endmember identification, as presented in Section 4.4. The task consists of the following subtasks:

T1.1 *Identify* candidate endmembers in the high-dimensional space

T1.2 *Compare* and *filter* candidate endmembers based on their spectral characteristics

T1.3 *Explore* the spatial distribution of pixels spectrally similar to a selected endmember

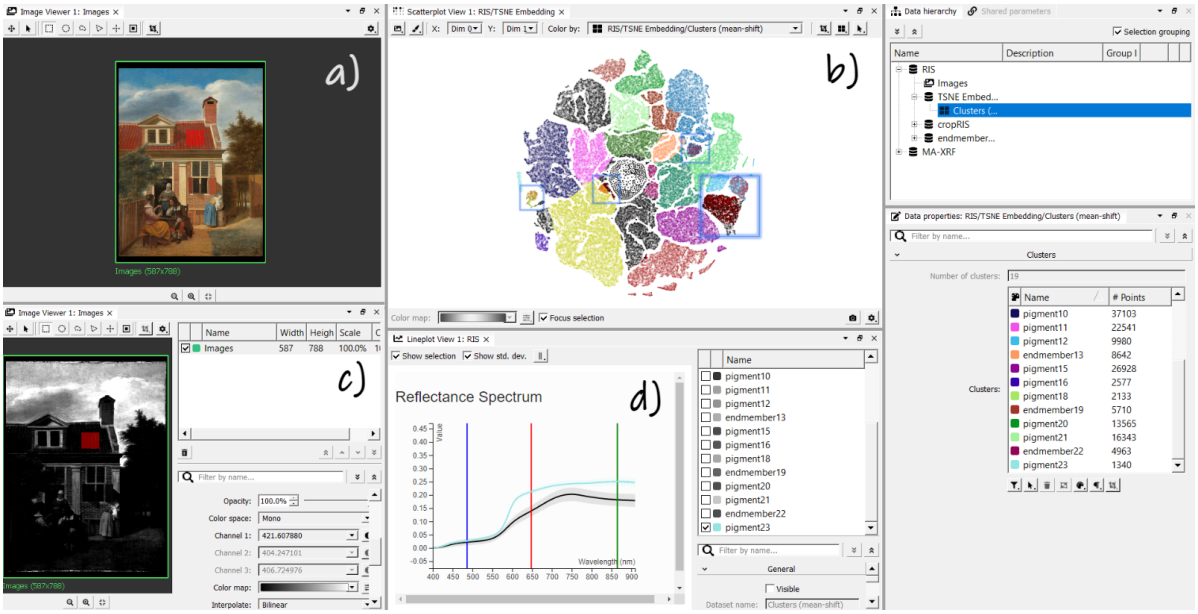

Figure 5.2: Workflow for endmember identification (**T1**); a) Selection of roof in photograph of painting in the Image View; b) Selection highlighted with blue boxes in the Embedding View; c) One wavelength of the RIS dataset in grayscale in the Image View; d) One endmember shown in blue and average spectrum of selection shown in black in the Spectrum View

The endmember identification task (**T1**) can be completed through the image, embedding and spectrum views (Figure 5.2). The histogram view can also be used as an addition to the image view. For this task we use the RIS data cube.

For identifying candidate endmembers (**T1.1**), we need to explore the spectra of the pixels. We can start by selecting pixels in the image view (Figure 5.2a) or in the embedding view (Figure 5.2b) and looking at the average spectrum of our selection in the spectrum view (Figure 5.2d). The average spectrum is shown in black in the spectrum view.

The image view (Figure 5.2a) allows us to select areas in the image which contain pixels that have a similar color. By selecting different areas in the image, we can probe the RIS data and get an initial idea of the spectra and therefore of the candidate endmembers present in the dataset. The image view comes with settings that the user can change such that the shown image fits in their analysis. A photograph of the painting can be directly visualized in this view (Figure 5.2a). However, we cannot show the entire spectral information of a pixel on an RGB-screen. Hence, it is possible to pick up to three dimensions of the data cube. We can set one wavelength as the image channel and show its values using false colors, by setting a colormap (Figure 5.2b). Alternatively, we can pick three wavelengths and map them to the red, green and blue channels of the image. This can be done by using a drop-down list in the image view or by dragging the R, G and B-lines in the spectrum view (Figure 5.2d). In Figure 5.3a, the RGB lines are set at the visible red, green and blue wavelengths and hence we obtain the image in Figure 5.3b. The user can drag the lines across the dimensions of the spectrum and gradually see how the image colors change. In Figure 5.3c, the shown spectrum was used as a reference for setting the red channel value at around 900 nm and we obtain the image in Figure 5.3d. Setting false colors in the RGB image allows us to identify pixels that are perceptually the same, but different in the infra-red range, and thus help us identify new candidate endmembers.

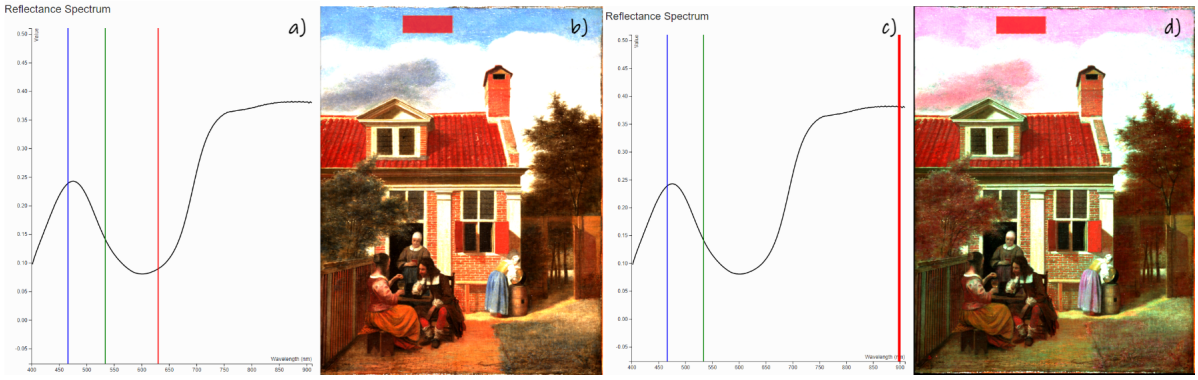

Figure 5.3: Examples of different RGB channels; a) Spectrum view with RGB lines set as the visible RGB wavelengths and b) Corresponding RGB image from Image View; c) Spectrum view with red line set in the near infra-red range and d) Corresponding RGB image from Image View

The image view uses a window/level linear function (Figure 5.4a) to set the dynamic range of the display colors of the pixels. In order to obtain accurate images, the user can set the window and level based on the data range, using a set of sliders in the image view. To make this more intuitive, we provide the histogram view, which shows one dimension of the dataset (Figure 5.4b) in order to more easily adjust the settings to the actual data. In the histogram we can visualize the window/level function in relation to the data. Hence, the user can interactively adjust the window and level by dragging the two ends of the diagonal line (highlighted in the blue boxes in Figure 5.4b) across the x axis and setting them to fit the data range. Additionally, the user can select one of the bars in the histogram and see the spatial distribution of the selected points in the image view.

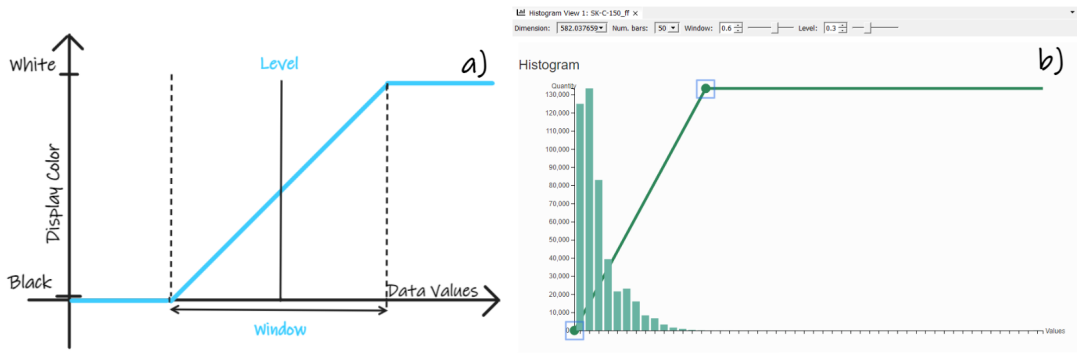

Figure 5.4: Window/Level function

Once we have an initial idea of the spectra in the image, we can select similar pixels in the embedding view (Figure 5.2b) and save these selections if their average spectra represent candidate endmembers. As discussed in Section 4.5, we cannot efficiently explore the data in the high-dimensional space directly. Instead, we use dimensionality reduction to project the data points into two dimensions and visually explore and interact with them. Since we can define candidate endmembers based on groups of pixels which have a similar spectrum, we chose the t-SNE algorithm [4] for dimensionality reduction. When computing the projection, t-SNE tries to preserve local neighborhoods between points in the high-dimensional space. We configured t-SNE to use the cosine distance for comparing points in the high-dimensional space. Hence, in our case, points which are spectrally similar appear close together in the embedding. Figure 5.5a shows the t-SNE embedding of the RIS dataset, where each point represents one pixel.

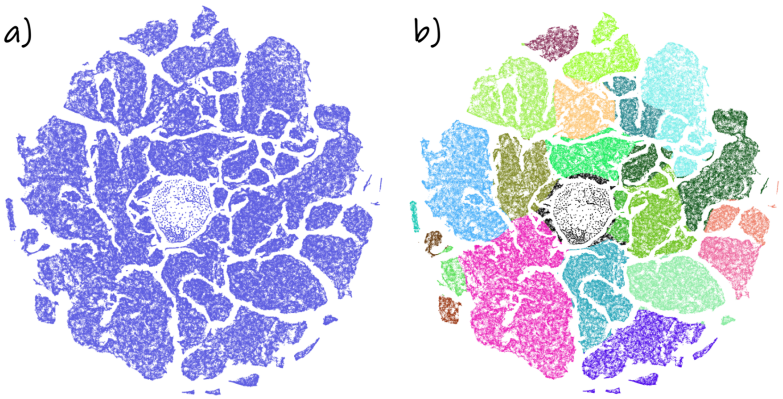

Figure 5.5: a) RIS t-SNE embedding; b) Clustered RIS t-SNE embedding (the clusters are indicated as the color overlay)

The user can then cluster the embedding and look at the average spectra of these clusters in the spectrum view in order to identify candidate endmembers. The clustering can be done manually, as the user can make selections in the embedding and save them as clusters or automatically, using the Mean Shift Clustering algorithm [63], as has been done in other applications [50]. Since the points are grouped based on their similarity, we want to make selections of neighbouring points, which result in average spectra with low standard deviation in the spectrum view. The standard deviation of an average spectrum is indicated as a semi-transparent area around it, as can be seen around the spectra in Figure 5.6b. In Figure 5.6b we can see a smaller standard deviation area for a selection of pixels similar in RGB, highlighted in pink in Figure 5.6a, compared to the standard deviation of a selection of more diverse pixels, highlighted in green in Figure 5.6a.

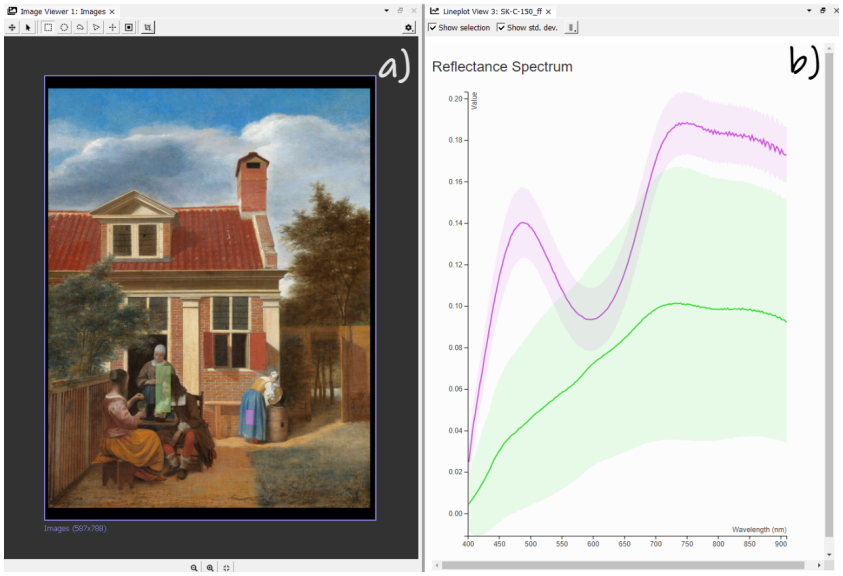

Figure 5.6: Spectrum standard deviation example; a) Two selections in Image view: selection of similar pixels in RGB shown in pink and selection of more diverse pixels shown in green; b) Indicated in the same colors, the corresponding average spectra of selections

Moreover, if we have a specific area in the painting for which we want to know the endmember, we can make a selection in the image view and based on it decide which points to cluster in the embedding. For example, in Figure 5.7a we selected part of the roof and we can see that in the embedding, Figure 5.7b, the selection appears mostly in the group of points surrounded by the blue box. Hence, we may decide to cluster that group and hence obtain a candidate endmember with a low standard

deviation representing the roof.

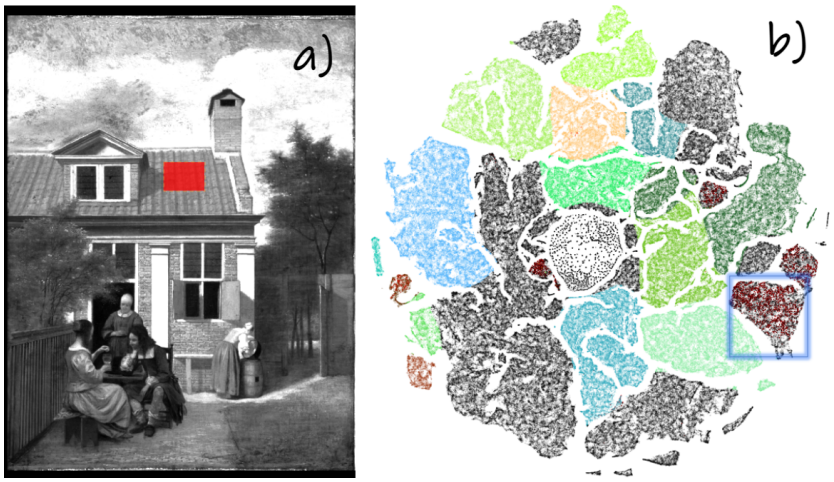

Figure 5.7: Clustering example; a) Selection of area in image; b) Part of selection highlighted in the blue box in the embedding, corresponding to one group of pixels

The purpose of the automatic clustering is to give the user a guideline for finding candidate endmembers. For example, the data can be over-clustered and then the user can manually merge clusters or discard them in order to get a final set of clusters corresponding to valid candidate endmembers with low standard deviations. We obtained the clustering in Figure 5.5b in this way. It represents just an illustration of the procedure, as we did not use any domain knowledge of realistic endmembers when refining the clusters.

We compare and filter the found candidate endmembers (**T1.2**) using the spectrum view. This view is implemented as a lineplot for which the x axis shows the wavelengths of the spectrum in nanometers (nm) and the y axis shows the measured reflectance. We can load all the candidate endmembers corresponding to the embedding clusters in this view. In Figure 5.8a, we can see a list of two endmembers on the right, out of which only the first one is visible. Based on only the endmembers, we can already assume the presence of some pigments. We can easily interact with the endmembers in the list and customize them by setting their name, color and visibility. By making some of the endmembers not visible, the user can focus on comparing only a subset of them. Moreover, the user can remove any duplicate or irrelevant endmembers during the analysis. In the end, we can export the filtered and visible endmembers to a file. Additionally, a set of predefined endmembers can be imported from a file in the view. For example, if the user has a set of endmembers from another similar painting of the same artist, they can load it into the view and compare it to the found endmembers.

We can explore the spatial distribution of pixels which are spectrally similar to an endmember using endmember and cluster maps (**T1.3**). Endmember maps are created in the spectrum view and visualized in the image view. We provide two different implementations for creating a map. One of them is the Spectral Angle Mapper, introduced in Section 4.3 and the other one is the Spectral Correlation Mapper [64]. We discuss their implementation in more detail in Section 5.3.2. With both algorithms, we obtain a grayscale map in which the pixels that are similar to an endmember are white and the pixels that are dissimilar are black. Additionally, we can map all pixels that do not meet a set threshold to black. We can choose to visualize a single endmember map in the image view (Figure 5.8b) in order to see how the map changes in real time as we adapt the similarity threshold in the spectrum view (Figure 5.8a). In order to decide the right threshold value, we use the MA-XRF data as well (Section 5.2). If we cannot find a map that accurately represents an endmember, we can still choose to discard the endmember (**T1.2**).

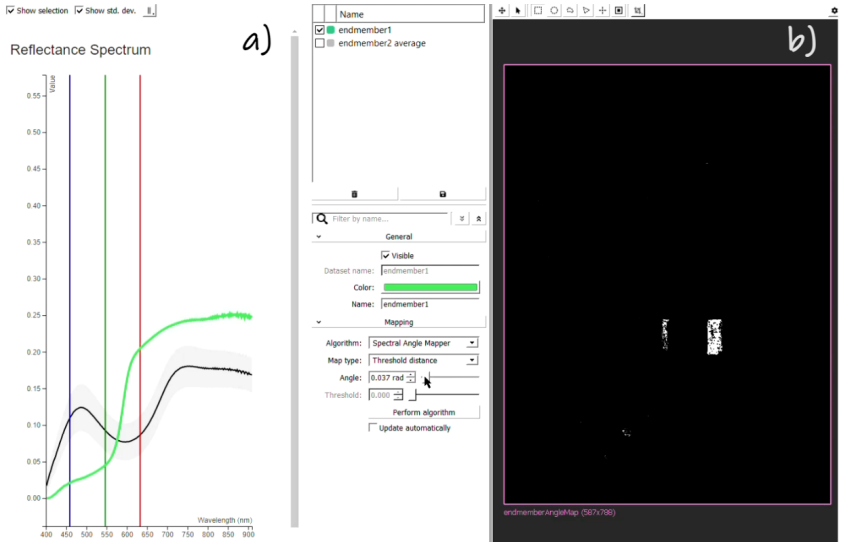

Figure 5.8: a) Spectrum View with one endmember shown; b) SAM map of shown endmember

Finally, we can load a combination of maps in the image view in order to visualize the spatial distribution of multiple endmembers at once and in order to potentially identify unexplored areas in the painting or areas that are missing endmember information (for example due to scanning errors). We can visualize multiple endmember maps by treating every map as a different semi-transparent layer (with a user selected constant colour) in the image, such as in Figure 5.9c,d. An alternative to the endmember map is the cluster map, which shows all the points belonging to a cluster. We can visualize the cluster maps in the image view, as shown in Figure 5.9a,b. Since we use clusters to define endmembers, the cluster map also shows all the points that are similar to an endmember. We further discuss the similarities between the endmember and cluster maps in Section 6.3.

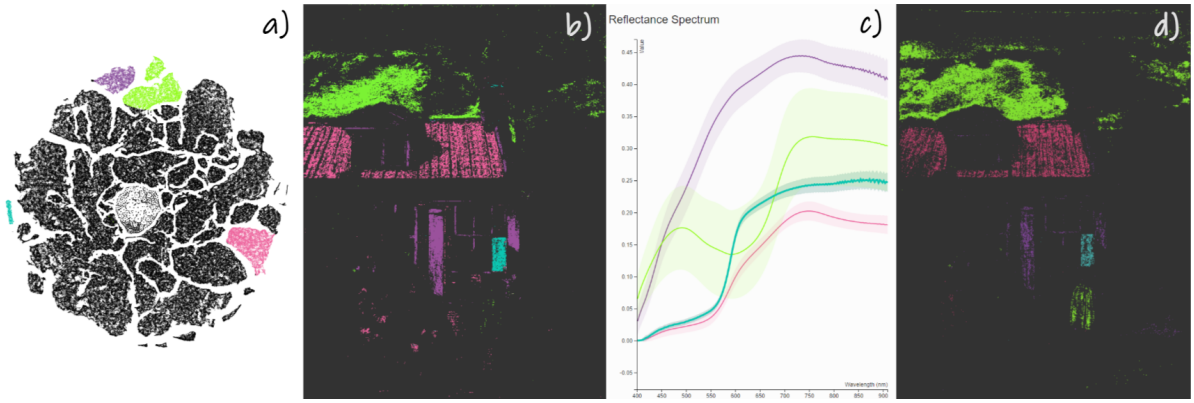

Figure 5.9: a) Four endmember clusters shown as different colors; b) Corresponding clusters in image view; c) Corresponding endmembers in spectrum view; d) Combined maps of the four endmembers

5.2. Pigment Identification

Here we present our design and implementation to tackle task **T2**, the pigment identification, as presented in Section 4.4. The task consists of the following subtasks:

T2.1 *Explore* the spatial distribution of pixels containing selected chemical elements

T2.2 *Identify* pixels that have similar molecular and elemental profiles

T2.3 For the pixels spectrally similar to an endmember, *identify* their elemental composition

T2.4 *Annotate* pixels according to the pigment they represent, based on both data modalities

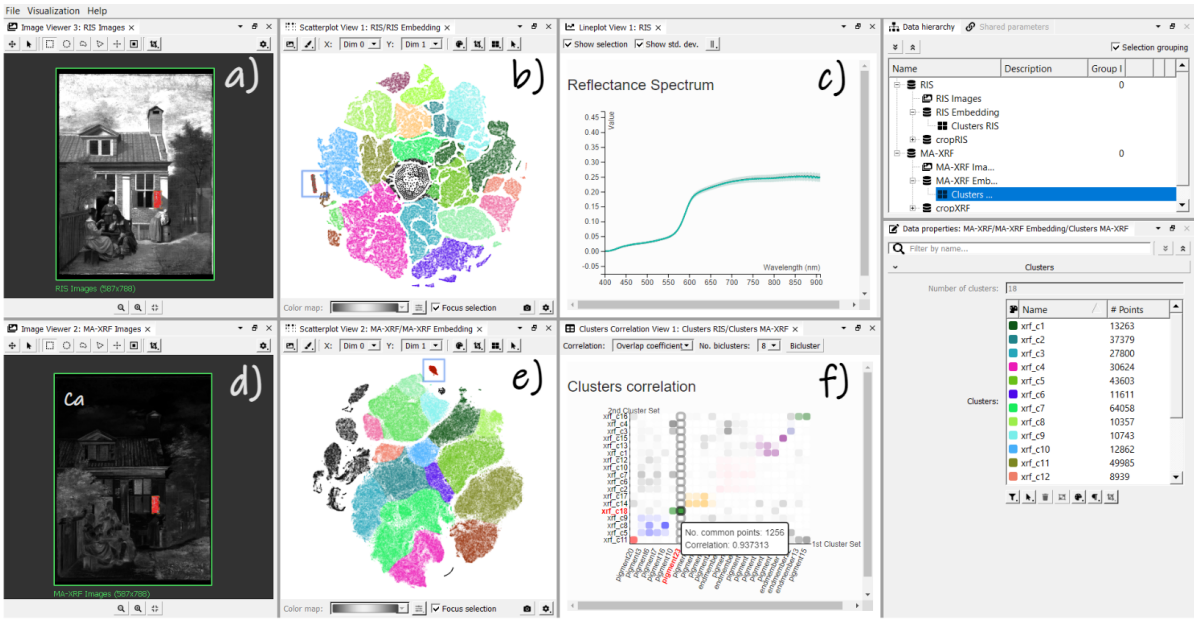

Figure 5.10: Workflow for pigmet identification (**T2**); a) RIS dataset in the Image View; b) Selected points highlighted by a blue box in the RIS Embedding View; c) Endmember highlighted on the x axis in f); d) MA-XRF elemental map in the Image View; e) Selected points highlighted by a blue box in the MA-XRF; f) Selection made in Cluster Correlation View

For the second task, the pigment identification (**T2**), we can use the image, embedding, spectrum and cluster correlation views (Figure 5.10). We want to characterise pixels with respect to both their molecular and elemental composition in order to identify pigments, so for this task we use both data modalities.

We can explore the spatial distribution of chemical elements (**T2.1**) by using the MA-XRF elemental maps. Similarly to visualizing the endmember maps (**T1.3**), we can use the image view in order to look at the elemental maps individually (Figure 5.11b) or in combination (Figure 5.11c). The combined map allows us to see the overlap between different chemical elements. For example, in case we have a mixture of pigments with distinctive elements in their composition and we see that those elements overlap, we can already assume that the mixture is present. Moreover, if a pigment is composed out of two elements, we could combine their maps and see where the respective pigment is present. Alternatively, we can use the embedding view in order to explore the overlap of two elements. In Figure 5.11a, we can see the selection of points that contain both Ca and Hg highlighted in the blue box. Based on the image in Figure 5.11b, we can see that the selection appears in an area where the Ca and Hg maps overlap in Figure 5.11c. In comparison to the image view, the scatterplot of two elements shows a clearer separation between the points that contain both elements and just one element. However, in the image view we can explore the overlap of more than just two elements. Hence, we can also use the elemental maps in order to help us decide the threshold angle in the SAM computation of the endmember maps (**T1.3**).

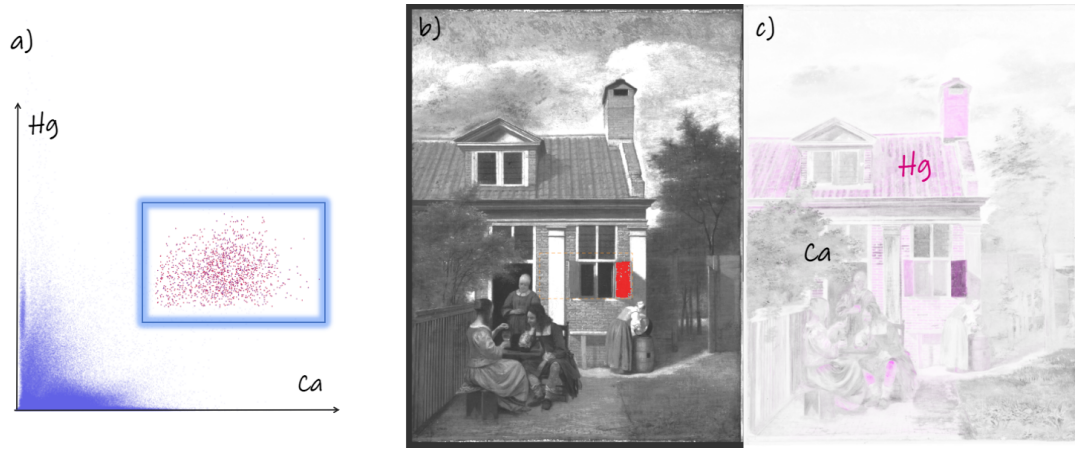

Figure 5.11: Example MA-XRF maps; a) Selection of points which contain calcium (Ca) and mercury (Hg); b) Ca elemental map; c) Combined elemental map of Ca (black) and Hg (pink)

In order to confirm any assumptions of pigments based on either one of the two data modalities, we need to analyse them in combination. Hence, we continue by exploring pixels with similar elemental and molecular compositions (**T2.2**). We expect that such points correspond to a pigment or a mixture of pigments. We already have clusters of points that have a similar molecular composition from the endmember identification workflow in Section 6.1.1, so we need to also find clusters of points with a similar elemental composition. If we have both sets of clusters, we can explore their overlap in the cluster correlation view (Figure 5.10f). Since the embedding shows us similar points, we can compute the t-SNE embedding of the MA-XRF data and cluster points with a similar elemental composition, just like we did with the RIS embedding. For the MA-XRF data we do not use the spectrum view, but we can use the endmember information as guidance for selecting the clusters. For example, we can open the RIS embedding and spectrum view (Figure 5.12b,c) in combination with the image view of the elemental maps and the MA-XRF embedding (Figure 5.12a,d). Then, we can make a selection in the embedding of the MA-XRF data, such as the one indicated in the blue box in Figure 5.12d and if it corresponds to an endmember in the RIS data we can save the cluster.

Figure 5.12: Clustering the MA-XRF data; a) Ca elemental map with selected points; b) RIS embedding with selected points; c) Spectrum of selection; d) Selection in MA-XRF embedding for creating a new cluster; e) Setting name of cluster in cluster list

After we obtain both sets of clusters (shown as the colors in Figure 5.14b,d), we can explore their correlation in the cluster correlation view (Figure 5.14a). This view is implemented as a heatmap and the color intensity of an entry corresponds to the similarity of the two clusters that intersect at that entry (in Figure 5.14a the similarity is measured using the overlap coefficient, described in Figure 5.14). A tooltip appears as we hover over one entry, showing the number of common points between two clusters and their similarity value. The entries are grouped based on their row and column similarity. The groups are indicated by the colors of the entries. Moreover, we can explore the spatial distribution of the common points in the image view (Figure 5.14c). Additionally, a cluster on any axis can be selected in order to further analyse it with respect to the other data modality. In Figure 5.14a we select the overlap of two clusters that are very similar (similarity of 0.9), as we can also see in the embedding views (Figure 5.14b,d), where the two corresponding clusters are highlighted in blue boxes. Hence, we identified a pigment at the intersection of the two clusters. Alternatively, we can load the clusters of one data modality as the color overlay in the embedding of the other modality and explore the overlap of clusters. For example, Figure 5.13 shows the RIS embedding colored with the MA-XRF clusters.

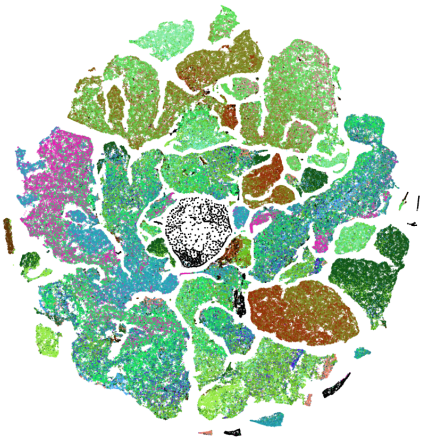

Figure 5.13: RIS embedding with color overlay of MA-XRF clusters

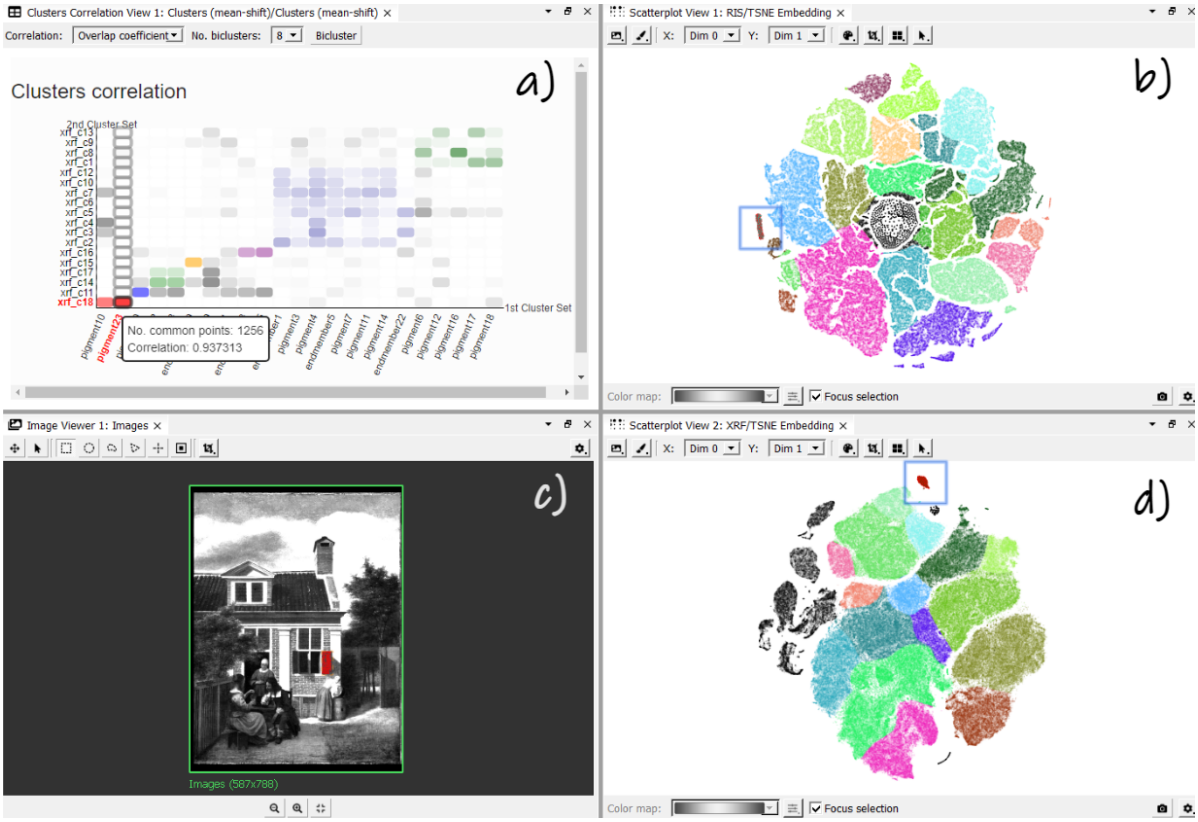

Figure 5.14: a) Selection made in Cluster Correlation View of RIS clusters (x axis) and MA-XRF clusters (y axis); c) Image View; Selection highlighted with blue box in b) RIS embedding and d) MA-XRF embedding

Since we now know the pixels with a similar molecular and elemental structure, we can find out what elements correspond to an endmember concretely (**T2.3**). For this, we can use the image view to visualize a combination of maps from both data modalities and observe their overlap. For example, if we assume the pixels in an endmember map represent a particular pigment, we can choose to combine this

endmember map with the elemental maps of elements that are present in the pigment in order to verify our assumption. In Figure 5.15b we can see an endmember map (corresponding to the endmember in Figure 5.15a) with two elemental maps (of Ca and Co) combined.

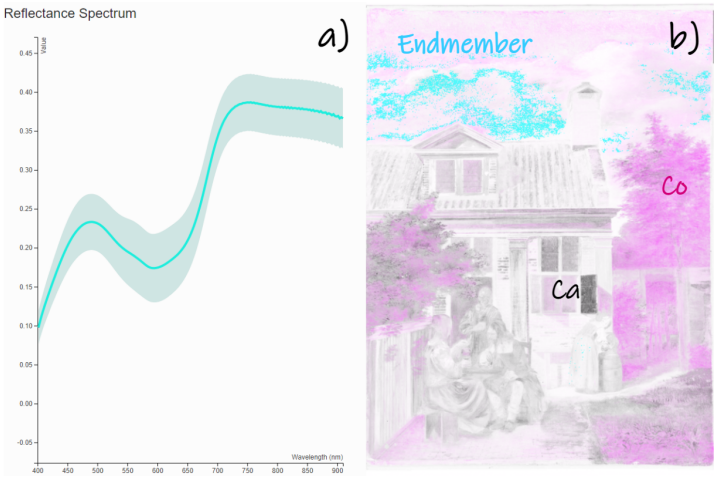

Figure 5.15: Combined map of both data modalities; a) Endmember in Spectrum View, b) Combination of endmember map in blue (corresponding to endmember in a)), Cobalt (Co) map in pink and Calcium (Ca) map in black

Moreover, we can pick one dimension of one modality as the color overlay in the embedding of the other modality. In Figure 5.16 we can see the RIS embedding with the distribution of two chemical elements (Hg and Co). In Figure 5.16a,b,c we highlight the points that belong to a cluster in the embedding and have a high value for the shown element (in the orange box the points with Hg and in the green box the points with Co). We know that these points are similar to an endmember and contain the shown chemical element.

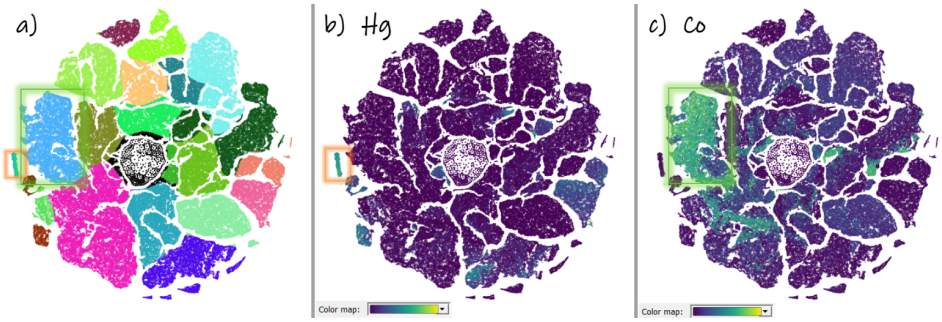

Figure 5.16: a) RIS embedding with color overlay of two elements, in a) Mercury (Hg) and in b) (Co); The green boxes highlight the points which have a high value for the chosen element

Similarly, if we consider the scatterplot of two elements discussed in Section 5.1 (Figure 5.11a), we can set the color overlay of the points as one or more of the endmember clusters and hence identify endmembers which contain one or both of the selected elements. For example, in Figure 5.17 we set the color overlay as the endmember clusters in the scatterplot of mercury (Hg) and calcium (Ca) and we can see that the group of points that contain high values for both elements (highlighted in the orange box) belong to one endmember cluster.

Lastly, for an overview of all endmembers and their elemental distribution, we can look at the correlation of the endmember maps and the elemental maps. The two map datasets can be clustered using the cluster dimensions plugin (Section 5.3.6) and visualized in the cluster correlation view (Figure 5.19b). The user can decide which points

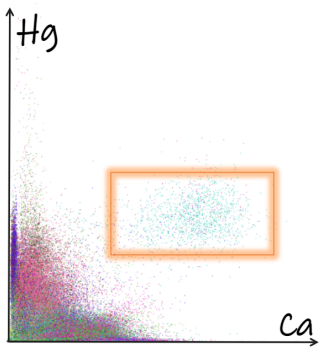

Figure 5.17: Scatterplot of Hg and Ca with endmember clusters set as color overlay with one endmember cluster that contains high values for both elements highlighted in the orange box

should be included in the cluster of every dimension. The clusters obtained in this way can overlap and they do not have to cover all the points in a dataset. In the case of the MA-XRF maps, it may be useful to refine the points considered to contain a chemical element as there can be noise in the elemental maps. We can interactively set a minimum threshold for every cluster, as shown in Figure 5.18d for the Ca map cluster. The Ca map is used as guidance both in the image view (Figure 5.18a) and in the embedding view as the color overlay (Figure 5.18c). The cluster points are shown in red in the embedding (orange boxes in (Figure 5.18b) and they correspond to areas that have a high value for Ca, as highlighted in Figure 5.18c with the same boxes. However, in the case of the RIS binary maps, we do not have to refine the clusters, as the default minimum threshold is set as the mean value of a dimension and hence only the pixels similar to an endmember are included in an endmember map cluster.

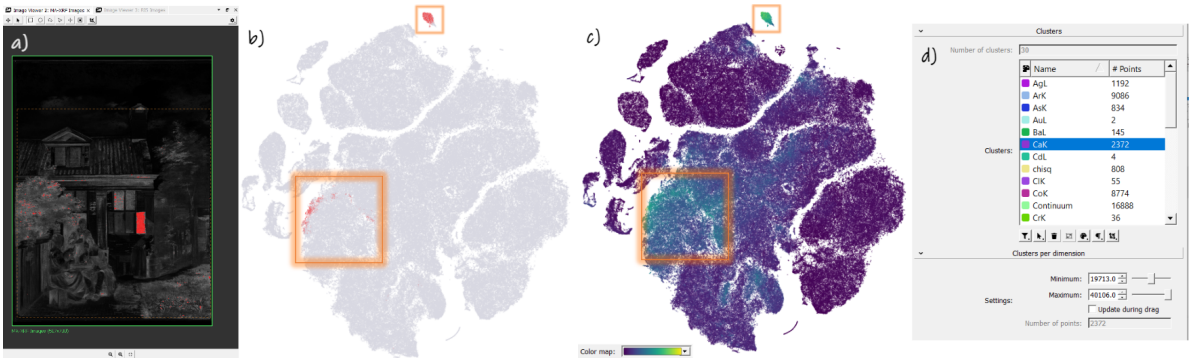

Figure 5.18: Adjusting cluster based on Ca elemental map; a) Ca elemental map with pixels included in the cluster highlighted; b) MA-XRF embedding showing the pixels included in the clusters in the two orange boxes; c) MA-XRF embedding with color overlay of Ca map showing pixels with high values for Ca in the orange boxes

If we use these two sets of clusters, one entry in the cluster correlation view represents the points that are both similar to a specific endmember and contain a specific chemical element. We can consider one endmember map cluster, indicated in Figure 5.19b as the highlighted column and go through every entry in the column in order to explore its elemental distribution.

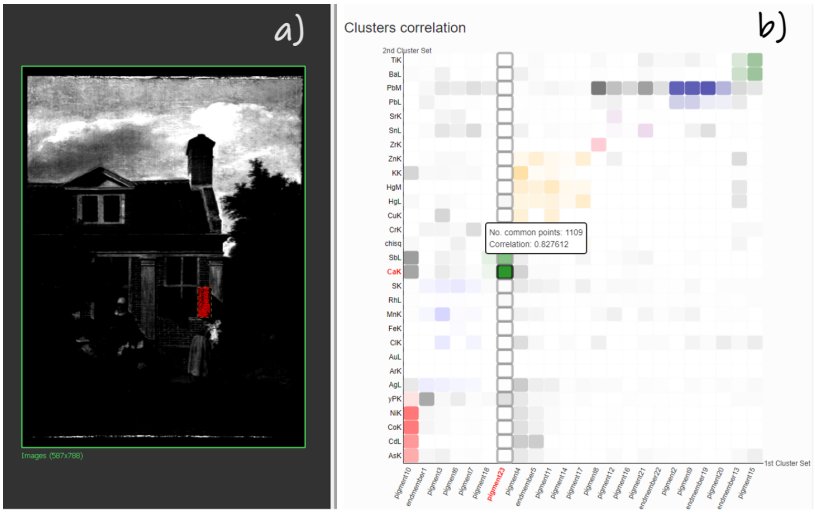

Figure 5.19: Exploring the elemental distribution of one endmember; a) Selection highlighted in the Image View; b) Cluster Correlation View of endmember (x axis) and elemental (y axis) map clusters with the focus on the endmember called *pigment23* and the selection of its intersection with the Ca map cluster

We can annotate a group of pixels according to the pigment they represent (**T2.4**) when naming the clusters in any modality. By giving pigment names to the clusters, we keep an organized list and can easily identify and filter clusters. Cluster names can be updated in the cluster list (Figure 5.12e) or, for the RIS data, also in the endmember list in the spectrum view (Figure 5.8a).

5.3. Implementation

The prototype is implemented in a plugin-based system, called HDPS, which is designed for the analysis of high-dimensional data. The system is implemented using Qt [65], which is a cross-platform user interface framework. The plugin-based nature of the system allowed for the easy extension of the existing functionalities. All the data loaders, the t-SNE and Mean Shift Clustering algorithms, and the views described in this chapter are implemented as plugins. The system already had an image loader, which we used for loading the photograph of the image and the MA-XRF elemental maps, and an ENVI loader for the RIS data cube. An endmember loader was added for our application for importing endmember lists. The image and embedding views were already present in the system. We added the spectrum, cluster correlation and histogram views. Additionally, we implemented two other plugins for processing the data, one for obtaining the complement of a Cluster dataset and one for clustering the dimensions of a Point dataset. The interface of our plugins remains consistent with the design of the system. Section 5.3.1 briefly describes general aspects of the system, which are relevant to our prototype. The remaining subsections describe the implementation of the five added plugins.

5.3.1. System Information

The image and embedding views are implemented using OpenGL [66]. The image view does not support the visualization of a zoomed in selection of one image, as shown in the design in Figure 4.6, but the user can open two different image views and set different zoom levels for them in order to obtain a similar result. Moreover, the system provides an improved version of t-SNE [67], which efficiently computes embeddings of millions of pixels.

The system also supports linking of parameters between views. The user can publish a parameter and link it to any other compatible field. We use this functionality for connecting dataset dimensions across views, like the RGB image channel values to three wavelength values in the spectrum view, as described in Section 5.1. This functionality was inspired by the needs of our application.

Moreover, multiple file types can be loaded into the system. The RIS data cube is loaded based on its header file (.hdr) using the endmember loader, for which we implemented nearest neighbour subsampling. Endmember lists are loaded using the endmember loader and saved as .txt files in the spectrum view. Through the image loader we can add images of various formats. For example, for the painting described in Section 4.2, its photograph is a .png file and the MA-XRF elemental maps are .TIFF images.

The three main data types of the system are Points, Images and Clusters. The RIS and MA-XRF datasets are loaded as Points and Images datasets. Points represent a group of high-dimensional pixels. Images extend Points by spatial coordinates. Clusters represent sets of pixels from a Points datasets. A user can save a selection of Points as a subset (a Points dataset) or as a cluster in a Clusters dataset. There is no requirement in the system for clusters to be complete (i.e. color all the data) or exclusive.

For our implementation, we used the JavaScript library D3 [68] for rendering the views, embedded in a Qt webview. The data for the views is processed using C++ such that VisuaLayered can run heavy computations efficiently. The communication between the backend and the frontend is handled by using QtWebEngine [69], which uses an embedded Chromium browser. Every plugin operates on or visualizes one or two datasets.

5.3.2. Spectrum View

The spectrum view, shown in Figure 5.20, is based on D3 line charts. The view uses one Points dataset as input. The average spectrum and standard deviation for a selection is obtained by computing the average of the selected pixels for every dimension of the dataset. This happens on the C++ side for performance reasons, as indicated previously. The boundaries of the shown standard deviation area are obtained by subtracting and adding the standard deviation to the average spectrum values. Additionally, a subset or a cluster set can optionally be added in the view to serve as endmember objects. A subset can be added in two ways, as one endmember per pixel or as one endmember for the average of the selection. When adding a set of clusters, each cluster is converted into one endmember. The properties of an endmember are: name, color, average spectrum data and source dataset. Any changes in the endmember cluster sets, like for example merging clusters, are reflected in the endmember list as well. Furthermore, a list of endmembers can be exported and reused in

a different project by importing. Finally, it is also possible to remove endmembers from the list. In comparison to the endmember selection design (Figure 4.7), the spectrum view is now a single line plot combining the RGB sliders, the average spectrum and endmember lines. The possibility to set an endmember's visibility and the fact that the typical number of endmembers is relatively small, allows us to show everything in one view.

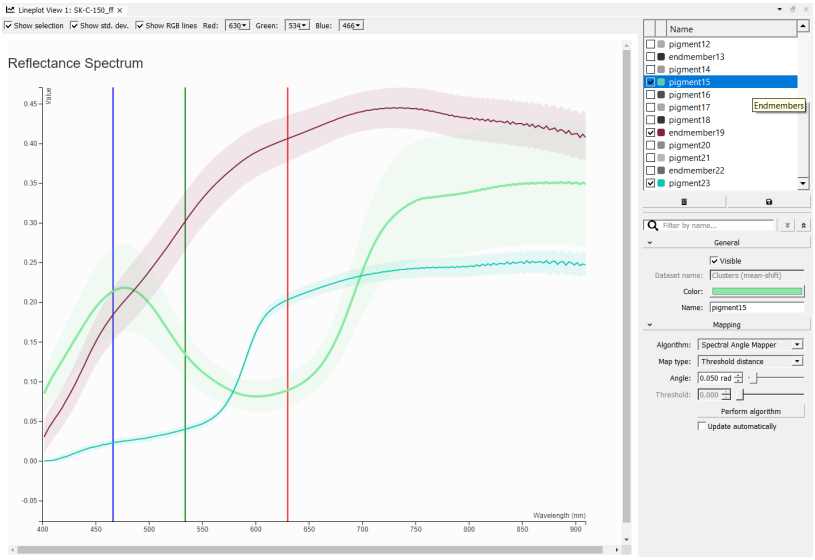

Figure 5.20: Spectrum View

For the endmember map computation, we provide two algorithms, the Spectral Angle Mapper and the Cluster Correlation Mapper. The Spectral Angle Mapper (SAM) [64], which is generally used in the field, is implemented using Equation 5.1,

$$\alpha = cos^{-1}\frac{\sum XY}{\sqrt{\sum(X)^2 \sum(Y)^2}}, \text{where} \tag{5.1}$$

α is the angle between a pixel spectrum and the endmember, X is the pixel spectrum, and Y is the endmember.

The Spectral Correlation Mapper (SCM) [64] is implemented using Equation 5.2,

$$Similarity = \frac{\sum (X-\overline{X})(Y-\overline{Y})}{\sqrt{\sum (X-\overline{X})^2 \sum (Y-\overline{Y})^2}}, \text{where} \tag{5.2}$$

X is the pixel spectrum, $\overline{X}$ is the mean of X, Y is the endmember, and $\overline{Y}$ is the mean of Y

Both algorithms consider the pixels as high-dimensional vectors. SAM uses the cosine distance in order to find similar pixels to an endmember, while SCM uses the Pearsonian Correlation Coefficient. For SAM, the similarity is the measured angle, which takes values in the range $[0°, 90°]$. A small value of α indicates high similarity and a large value indicates low similarity. Figure 5.21 shows an example of the measured angle in two dimensions. For SCM, the similarity is a value in the range $[-1, 1]$, where a value of 1 indicates a strong positive correlation and a value of -1 indicates a strong negative correlation. SCM is an improvement of SAM as it shows negative correlations in the data, while SAM only considers absolute values [64].

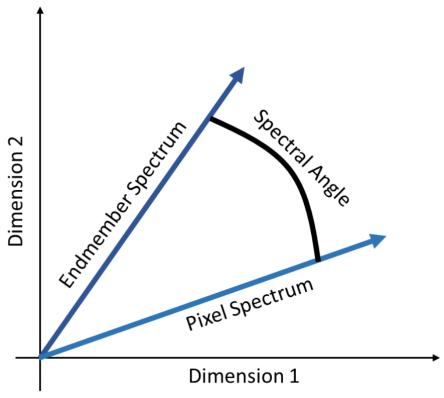

Figure 5.21: SAM illustration in two dimensions

Every computed endmember map is saved as a dimension in the endmember map dataset. We create one such

dataset per mapping algorithm such that the obtained maps can easily be compared. We provide three types of maps. The first one is binary, which assigns 1 to the pixels that are similar to an endmember and 0 to the other pixels. The second one is the threshold based one, which shows the similar pixels to an endmember based on the set threshold (for SAM, it shows the pixels at distances lower than the threshold; for SCM it shows pixels which have a similarity value higher than the threshold). The third one is distance based, showing all the distances between the points in the data and an endmember. Only the distance based one is saved in memory and the first two ones are computed on the fly based on the distance map, whenever the threshold angle is changed. In comparison to the design of the mapping analysis (Figure 4.8), a combined map is not only one single image, but can be obtained through the layering of multiple images, for which we can set the visibility, which gives us more flexibility in the combinations we can obtain.

5.3.3. Cluster Correlation View

The cluster correlation view, shown in Figure 5.22, is based on D3 heatmaps. The view uses two Clusters datasets as input. One entry in the heatmap shows the similarity between two clusters, based on the number of points they have in common. The similarity is computed on the C++ side for all possible combinations of two clusters and sent to the JavaScript side. Clicking on one of the entries of the heatmap or on one of the cluster names on the axes, sends a signal to the C++ side to select the corresponding pixel indices in all of the other views.

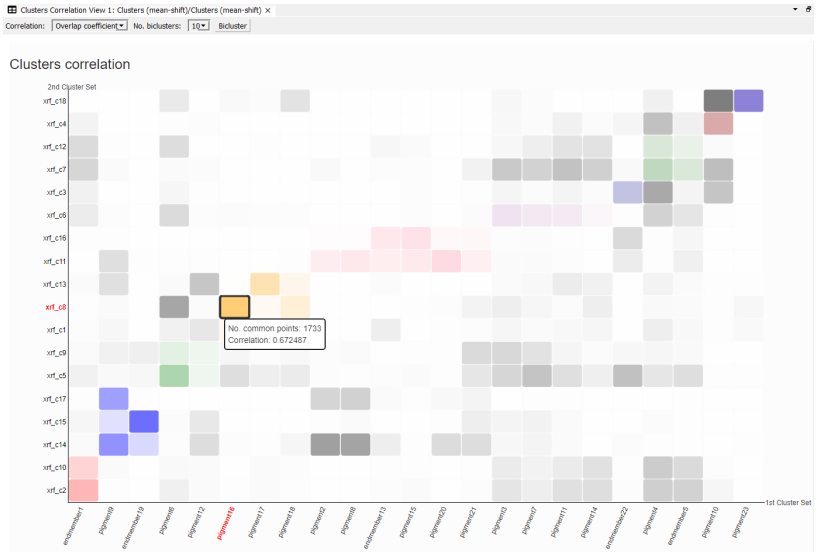

Figure 5.22: Cluster Correlation View

The correlation value, shown in the heatmap as the intensity of the entries' colors, can be computed using one of the following similarity measures, which have values in the range $[0, 1]$:

1. Overlap coefficient [70]: $\text{Overlap}(X, Y) = \dfrac{|X \cap Y|}{\min(|X|, |Y|)}$
2. Jaccard index [71]: $JI(X, Y) = \dfrac{|X \cap Y|}{|X \cup Y|} = \dfrac{|X \cap Y|}{|X| + |Y| - |X \cap Y|}$
3. Sørensen–Dice coefficient [72]: $SDC(X, Y) = \dfrac{2|X \cap Y|}{|X| + |Y|}$

, where X, Y are two finite sets.

Moreover, a tooltip appears when the user hovers over an entry in the heatmap showing the number of common points between two intersecting clusters and their similarity coefficient. While all three

measures show the similarity between two finite sets, they provide slightly different results. For example, if one set is a subset of the other one, the overlap coefficient results in a strong correlation, a value of 1, while the other two metrics result in lower correlation values. By default, the overlap coefficient is selected. The user can choose any of the three measures.

The entries in the heatmap are ordered using the Spectral Co-clustering algorithm [73]. In the last step of the algorithm, the heatmap entries are biclustered with k-means, using the *Clusterfck* library [74]. Entries that have similar rows and columns are clustered together. The colors used in the heatmap indicate the clustering and gray entries do not correspond to any cluster. Furthermore, the user can set the number of biclusters.

5.3.4. Histogram View

The histogram, shown in Figure 5.23, is based on D3 barplots. The view uses a Points dataset as input. Per dimension, the data points are sorted in a given number of intervals on the C++ side. The number of points in every interval is sent to JavaScript. The histogram can be used for setting the window/level settings of the image view, as described in Section 5.1. The user can set the minimum and maximum values of the window based on the data values in the histogram by moving the two ends of the diagonal line, which are represented by two circles. The level is computed automatically as the middle point of the window.

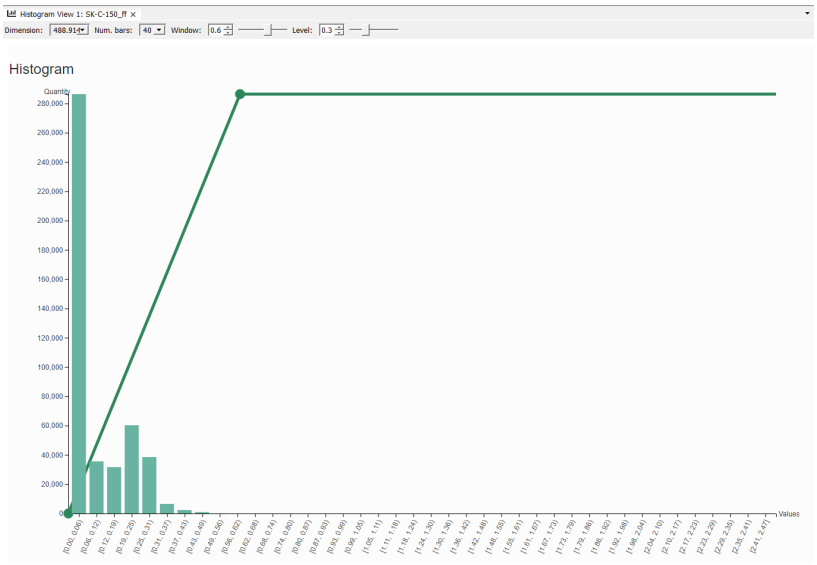

Figure 5.23: Histogram View

5.3.5. Complement Cluster

This plugin is fully implemented in C++. Given a Clusters dataset and its source, which is a Point dataset, this plugin identifies all the points that are not included in any of the clusters, but are part of the source dataset. The resulting points are saved as a new cluster in the input Cluster dataset. The user can use this complement cluster in order to identify areas in the painting that are unexplored.

5.3.6. Cluster Dimensions

This plugin takes as input a Points dataset and is fully implemented on the C++ side. It creates a cluster for every dimension of the dataset. The points that are initially included in the cluster have a value that is larger than or equal to the mean value of a dimension. The user can adapt the minimum and maximum range of the cluster. This plugin turns a Point dataset into a Cluster dataset. This is useful, for example, for converting the endmember maps into Cluster sets for comparison and further processing.

6

Evaluation

We conducted an evaluation with our collaborator from the Rijksmuseum in order to assess the usability and effectiveness of VisuaLayered. The evaluation consists of three parts. In the first part, we asked our collaborator to complete five painting analysis tasks using her typical analysis workflow, for which she mainly uses the geospatial analysis software, ENVI. In the second part, we asked her to complete the same five tasks using VisuaLayered. She analysed the same painting using both workflows. We observed and made an audio recording of her workflow in both systems and recorded her screen for the analysis in VisuaLayered. Based on all of these materials, we describe and compare the two workflows. In the third part, we asked her to fill in a survey via Microsoft Forms [75], which is aimed at assessing the features and usability of VisuaLayered and at comparing the two workflows. The survey can be found in Appendix A.2 and consists of 6 main sections. In the first section we ask the profession of our participants. The second section is the list of tasks to be completed using ENVI. The third section is the same list of tasks for VisuaLayered and each one or two tasks is/are followed by corresponding questions answered on a 5-point Likert scale [76]. The fourth section consists of questions about the comparison of the two systems. The fifth section asks five general questions about VisuaLayered. The last section consists of 10 standard System Usability Scale (SUS) questions [77].

The evaluation lasted one day. We met with our collaborator at the Rijksmuseum and we conducted the evaluation in her normal work environment, using her work PC. For each of the two workflows, she spent $\approx$ 2-3 hours. In general, she spends more time in order to analyse a painting, but since we wanted to compare both workflows back-to-back in one day, we decided to split the available time equally between the two workflows and see how much she can achieve, without rushing through the tasks. Our collaborator was involved in the development process from the beginning, so at the moment of the evaluation, they were already familiar with parts of the system. However, we still provided a written example workflow with screenshots of possible ways of combining the available views. During the evaluation we answered any questions regarding these provided examples.

In Section 6.1 we describe the workflow of our collaborator in both systems. Then, we go over the survey results in Section 6.2. Lastly, we discuss additional results obtained using VisuaLayered, that are not part of the case study workflow and the limitations of the evaluation in Section 6.3.

6.1. Case Study

In order to prepare for the evaluation, we asked our collaborator to provide two datasets, that they did not analyse in the past, one for RIS and one for MA-XRF, of the same size for a painting of her choice. They provided data on the painting *Still Life with a Vase of Flowers and a Dead Frog* by Jacob Marrel, 1634. Jacob Marrel is a German still life painter who lived in the 17th century and spent part of his life in Utrecht, The Netherlands. The painting is oil on panel and has a size of 31cm x 40.3cm (w x h). The MA-XRF data was acquired and processed by Nouchka De Keyser and the RIS data was acquired and stitched by Francesca Gabrieli. Both datasets have 1257 x 1635 pixels (after re-scaling). The RIS data cube has 256 dimensions and for MA-XRF we have 34 elemental maps. We prepared the t-SNE embeddings for both data modalities in VisuaLayered in advance of the evaluation.

As described above, we prepared 5 painting analysis tasks to be completed by our collaborator. We

go over them one by one and explain their connection to the task abstraction provided in Section 4.4. We started the evaluation with the typical analysis workflow in ENVI, with the goal to create results for comparison, then we moved on to VisuaLayered.

In Section 6.1.1 we describe the first two evaluation tasks which refer to task **T1**, the endmember identification. In Section 6.1.2 we describe the remaining 3 evaluation tasks which refer to task **T2**, the pigment identification. For all the images showing views from VisuaLayered, selected pixels and points are indicated in red.

6.1.1. Endmember Identification

We wanted to evaluate how well the user can identify endmembers in VisuaLayered with the first two evaluation tasks:

- **ET1** Identify the endmembers present in the studied painting (**T1.1**, **T1.2**).
- **ET2** Explore the spatial distribution of pixels that are spectrally similar to the found endmembers (**T1.3**).

ENVI For **ET1**, our collaborator started with opening the RIS dataset in an image viewer (Figure 6.1) in ENVI. By selecting different areas in the RGB image (Figure 6.1a) with the help of two zoomed in views (Figure 6.1b,c) and looking at the selected pixels' spectra in a spectrum view (Figure 6.1d), she gained an initial idea of the present endmembers. Then, with a false color image (Figure 6.2), she observed the underlying and otherwise hidden pigment used for indicating the position of three flowers, which is a technique commonly used by the painter. The hidden pigment appears in red in the false color image, as the red channel is set in the near infra-red range and we can see it in Figure 6.2 in the three flowers that are highlighted in green boxes. After this, she started defining endmembers using two automatic methods and by manual selection in ENVI. The Sequential Maximum Angle Convex Cone (SMACC) algorithm automatically identified 30 endmembers. This was followed by using the Spectral Hourglass Wizard (SHW) which found 35 endmembers automatically. In the 3D reduced space obtained with the SHW, she also made manual selections of clustered pixels, looked at their spectra and found 7 endmembers (Figure 6.3). Identifying endmembers in this space was tedious as she had to rotate the space and try different combinations of principal components for the three dimensions in order to better understand the structure of the data. Afterwards, she went through the automatically found endmembers one by one in the spectrum view in order to verify whether their shapes are physically valid. In the end, she saved 11 endmembers from the SHW automatic computation, four from SMACC and all the manually identified ones, resulting in a final list of 22 endmembers as input for **ET2**.

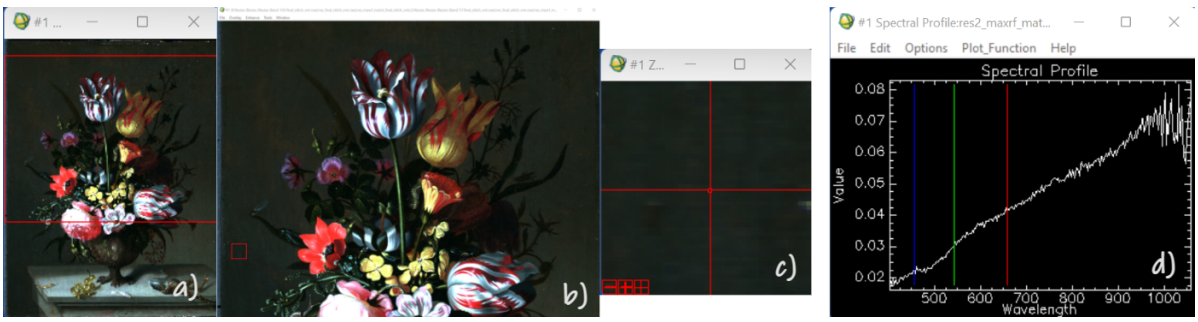

Figure 6.1: Visual inspection of the spectra in ENVI; a), b), c) Image view with different zoom levels; d) Spectrum view showing spectrum of selected pixel

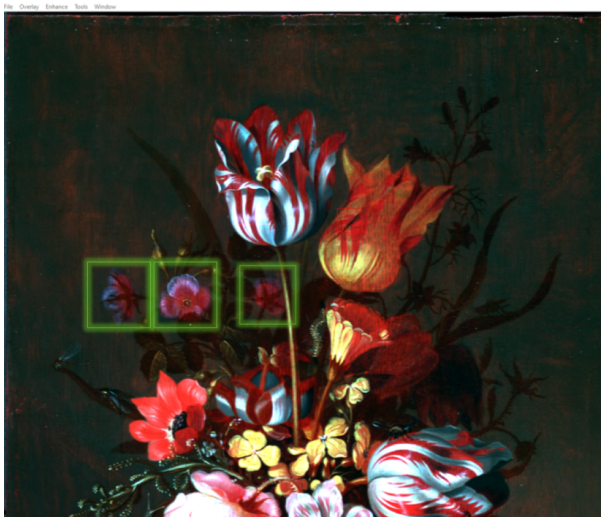

Figure 6.2: False color image in ENVI with the red channel in the near infra-red range showing hidden pigment in the three green boxes

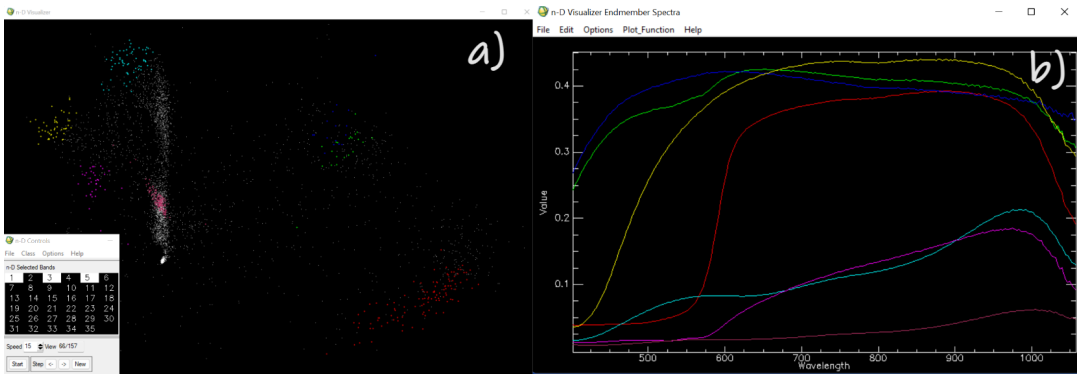

Figure 6.3: Manual endmember selection in ENVI; a) Clustered SHW embedding; b) Endmembers corresponding to the shown clusters

In **ET2** the goal was to explore the spatial distribution of pixels similar to the found endmembers. She did this by computing grayscale endmember maps using SAM (Section 4.3), checking their quality and saving eight maps (one for an automatically found endmember and seven for all the manually selected endmembers). As the decision of which endmember maps to keep is partially based on the MA-XRF elemental maps, we describe this part of the analysis in Section 6.1.2.

VisuaLayered With task **ET1** we wanted to observe how our collaborator identifies candidate endmembers and compares and filters them based on their spectral characteristics (**T1.1**, **T1.2**) in VisuaLayered. She started by opening the image, spectrum and embedding views (Figure 6.4b). Based on the structure obtained in the t-SNE embedding, she manually annotated groups of points in the data and assigned them to a cluster set. The clusters are indicated as different colors in the embedding. She selected groups of neighbouring points and looked at their average spectrum and spatial distribution in the image. If she found a relevant spectrum, she saved its corresponding points as a cluster and loaded it in the spectrum view as an endmember. Since t-SNE projects the data in 2D, in comparison to the 3D embeddings obtained in ENVI (based on different principal components), it was easier to identify groups of points for clustering in VisuaLayered. Furthermore, in comparison to ENVI, where she picked endmembers only based on the shape of the spectrum, in VisuaLayered she also assessed the quality of a candidate endmember based on its standard deviation. For example, in Figure 6.5 we can see a selection of points (Figure 6.5a) that was adjusted such that its endmember has a small standard deviation (Figure 6.5b). Additionally, in VisuaLayered she could make selections in the image

view in order to identify other relevant clusters in the embedding. In Figure 6.4a we can see the image selection mostly appears in three groups of the embedding, which resulted in creating the highlighted cluster in Figure 6.4b. In total, she manually created 19 clusters and therefore 19 endmembers.

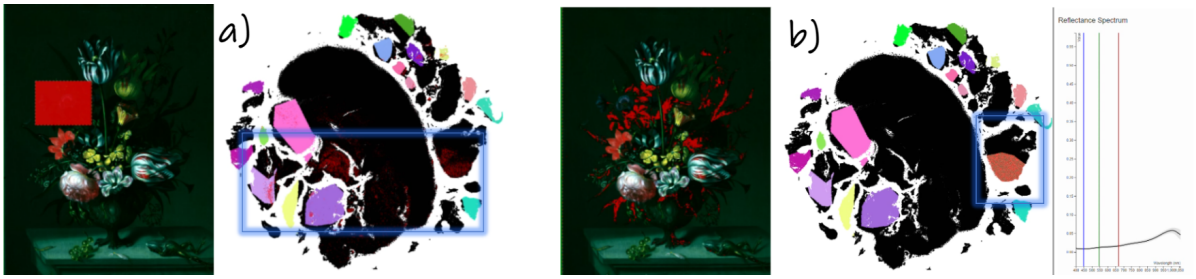

Figure 6.4: a) Selection in image for identifying clusters in the Embedding View; b) Highlighted cluster in the Embedding View created based on selection

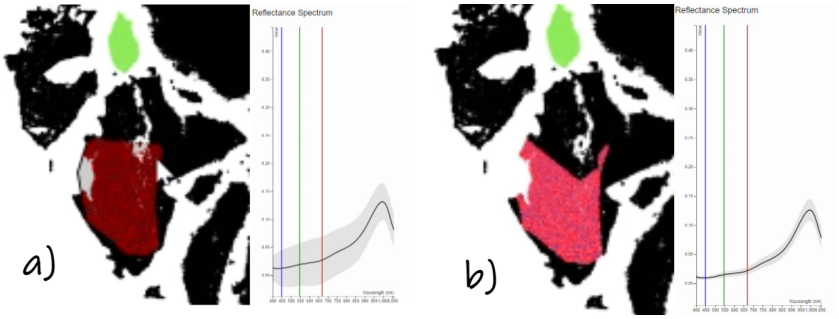

Figure 6.5: Use of standard deviation area for cluster decision; a) Initial selection; b) Adjusted selection that was saved as a cluster

Moving on to **ET2**, which is based on **T1.3**, she used SAM, which is available in the spectrum view, in order to create binary endmember maps. Setting the threshold angle for SAM was more difficult in VisuaLayered than in ENVI, as in ENVI the angle can be set with the help of a histogram. However, setting the visibility of endmembers in the spectrum view helped our collaborator keep the view organised. In ENVI the spectrum view sometimes became too cluttered with spectra and she had to open a new view in order to continue visualising endmembers. In VisuaLayered she only used one such view. In total, she saved 12 maps, from which six were also found in ENVI and six were obtained for endmembers identified only in VisuaLayered. However, there were also two endmembers only identified in ENVI. Figure 6.6 shows one of the six new endmembers from its cluster selection to its mapping. It represents the background of the image. For these results, we need to take into consideration the fact that she spent a shorter amount of time than she normally does in order to analyse a painting and identify endmembers. Hence, these results may not be representative of the number of endmembers found in the usually allocated time.

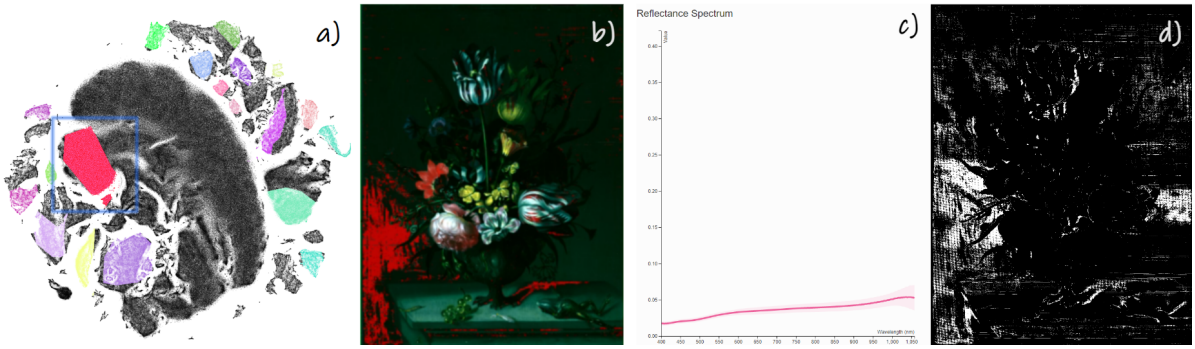

Figure 6.6: Background endmember; Endmember cluster selected in a) Embedding View, b) Image View; c) Endmember shown in Spectrum View; d) Obtained map for background endmember

6.1.2. Pigment Identification

Next, we wanted to evaluate how well the user can identify pigments in VisuaLayered with the following three evaluation tasks:

- **ET3** Consider the endmembers from task 1 (and the pixels that are similar to them) and establish their elemental composition as well, in order to identify the used pigments (**T2.1**, **T2.2**, **T2.3**).
- **ET4** Are you satisfied with your endmembers? If not, consider refining your endmember selections.
- **ET5** Can you reason about the possible order of the paint layers? If yes, annotate your selections to indicate this (name clusters e.g. top_ultramarine, bottom_ochre) (**T2.4**).

ENVI As part of her typical analysis, our collaborator used the MA-XRF elemental maps and RIS endmember maps in order to identify pigments for **ET3**. She started by creating an endmember map in ENVI with a base threshold value. Then, she looked at the spectra of the similar points highlighted in the map in order to check whether the threshold needs to be changed. If points with dissimilar spectra were considered similar to the endmember, she reduced the threshold and checked again. Moreover, she visualized the elemental maps using different software in order to validate the endmember maps. Based on an endmember, she already assumed a pigment and therefore searched for the elemental maps of elements that exist in the composition of that pigment. She visualized the maps as separate images and visually compared them. She tried to overlap the highlighted points in the endmember map with the highlighted points in the corresponding elemental maps as much as possible. After she computed all the maps, she was happy with the identified pigments and did not refine the endmember list for task **ET4**.

VisuaLayered In VisuaLayered, for task **ET3**, she used a combination of all the available views, except for the histogram view. As part of the endmember map creation process she used the embedding view to better understand the relation between the two data modalities (**T2.2**). By opening the embedding of the MA-XRF data and using one endmember map as the color overlay, she could observe the distribution of the points similar to an endmember in the MA-XRF data. For example, in Figure 6.7b the *end2* map is set as the color overlay in the embedding and its points appear mainly in two groups of the MA-XRF data, that are highlighted in the green boxes. Since she already assumed that the endmember represents azurite, which contains copper, she expected to also see the highlighted points in the copper map. Hence, she then changed the color overlay as the copper (Cu) map and instantly saw the same two regions highlighted, as shown in the green boxes in Figure 6.7c. However, the fact that the points appeared in two different groups indicated to our collaborator that one of the groups should contain an additional chemical element, which is not present in the other group. She tried a few more elemental maps as the color overlay until she settled on cobalt (Co), which is present only in the top cluster as shown in the green box in Figure 6.7d. Hence, this view offered a clear separation of points that have a similar reflectance spectrum, but a different elemental composition.

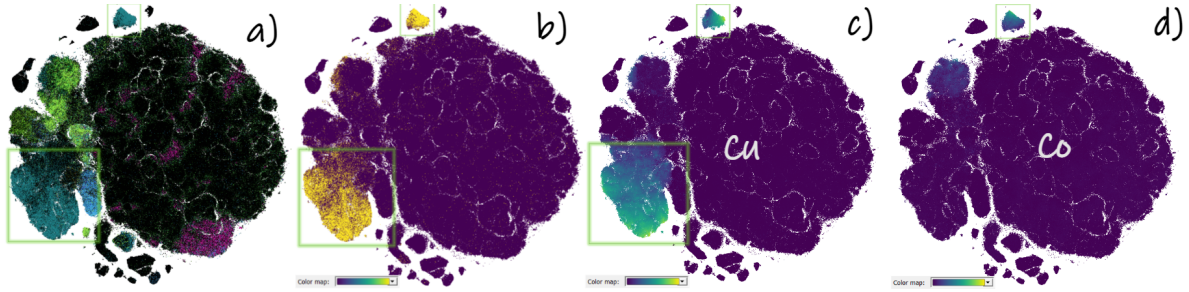

Figure 6.7: MA-XRF embedding with different color overlays: a) clusters of endmember maps, b) azurite endmember map, c) Cu elemental map, d) Co elemental map

Next, she wanted to establish the elemental composition of all endmembers (**T2.2**, **T2.3**) by analysing the endmember and elemental maps in combination in the cluster correlation view. For this, she automatically created one cluster per endmember maps, where each cluster contains all the points that are similar to an endmember. She also automatically created one chemical element cluster per elemental map (**T2.1**) and manually refined these clusters. In Figure 6.8a,b we can see the mercury (Hg) cluster as the selected points, with the minimum threshold value set by our collaborator in Figure 6.8c. The corresponding mercury map is used in the image view as the background (Figure 6.8a) and in the embedding view as the color overlay (Figure 6.8b) as guidance for setting the threshold.

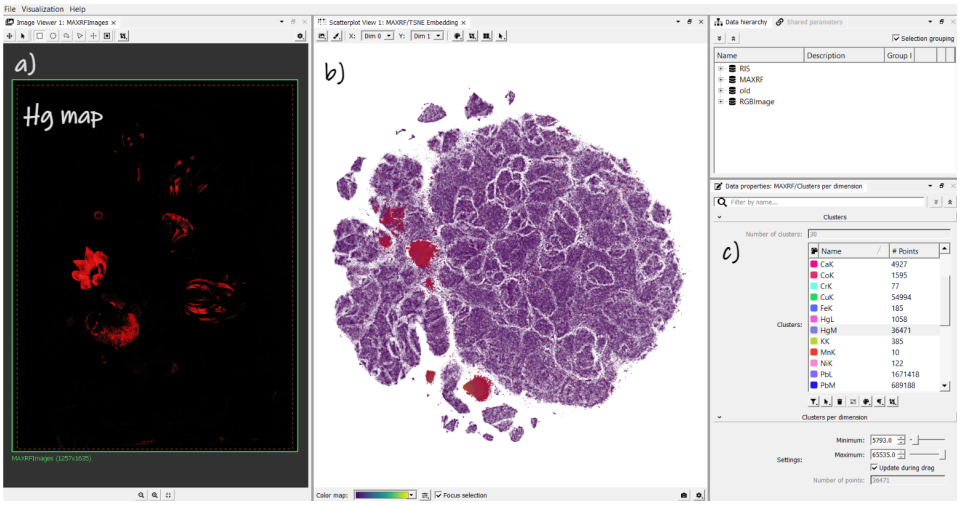

Figure 6.8: Adjustment of the Hg cluster; a) Cluster highlighted in red in Image View and b) Embedding View; c) Settings for cluster adjustment

She then loaded these two cluster sets in the cluster correlation view and saw the elemental composition of all identified endmembers (**T2.3**). Hence, this view offered a clear overview of the connections between the two data modalities. In the case of *end2*, it clearly showed the distribution of its map points in three main elemental clusters (Figure 6.9). Our collaborator could observe this separation in ENVI as well. However, in ENVI it was a process that was more time-consuming and more prone to error as she had to search for the right elemental maps and then visually inspect them in order to establish the overlap in her mind. With the cluster correlation view, this information was shown directly. Moreover, the image view was no longer used for establishing the overlap of the maps (like in ENVI), but for spatially visualizing the separation of one endmember cluster into its elemental clusters (Figure 6.9b,c,d). Moreover, the cluster correlation view provided new information on the overlap of maps that was otherwise missed in ENVI (e.g. *endmember 1* with the calcium (Ca) and potassium (K) maps).

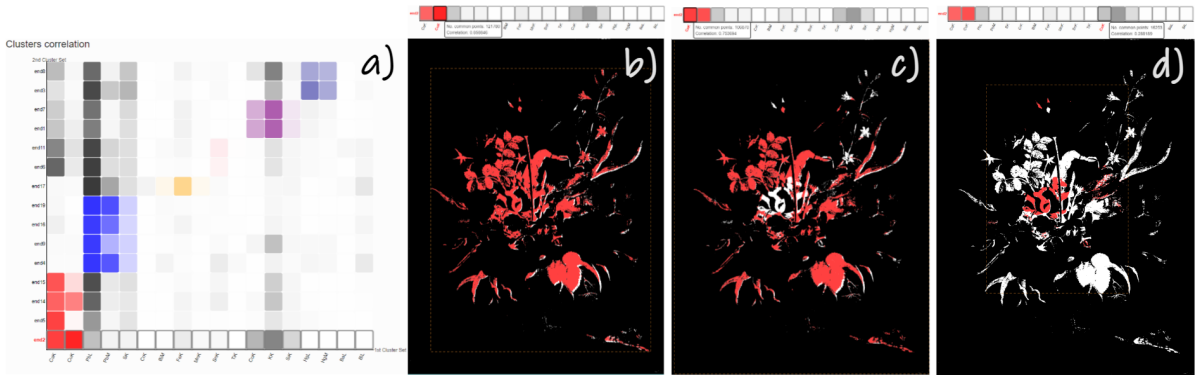

Figure 6.9: Example separation of endmember cluster into elemental clusters; a) Cluster Correlation View with *end2* map cluster highlighted; The background binary image in b), c) and d) is the *end2* map; b) Selection of Cu map cluster; c) Selection of Ca map cluster; d) Selection of Co map cluster

Our collaborator pointed out that the leaf in the azurite map obtained in ENVI (leaf indicated in the blue box in Figure 6.10a) is mapped better in VisuaLayered (leaf indicated in the blue box in Figure 6.10b). She knew this because azurite contains copper (Cu), which is present in the leaf as shown in the blue box the MA-XRF copper map (Figure 6.10c). When comparing the two endmembers for azurite (Figure 6.10d), we can see that the VisuaLayered endmember has a higher peak at around 980 nm. The endmember in VisuaLayered is obtained based on a cluster of points in the embedding which mostly contains points from the leaf. Hence, this endmember can more accurately map the leaf, compared to the ENVI one.

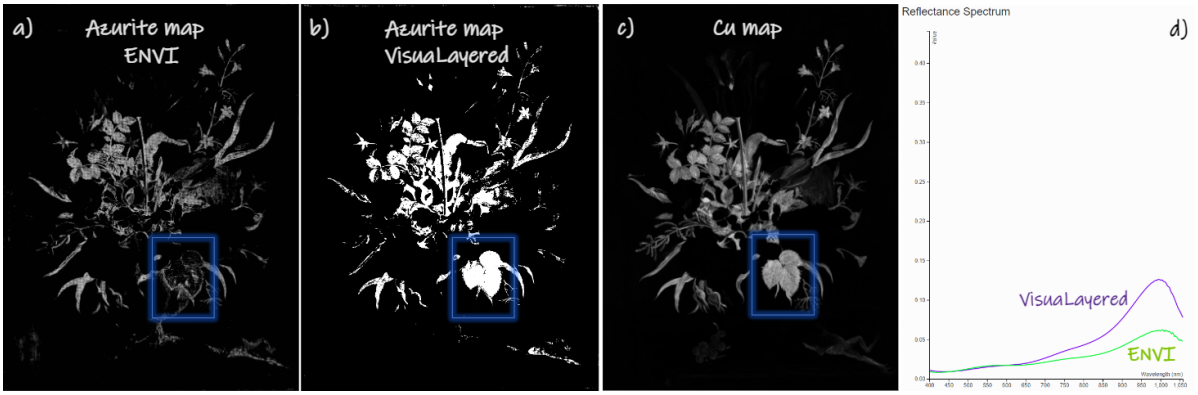

Figure 6.10: Azurite map comparison; a) ENVI endmember map; b) VisuaLayered endmember map; c) Cu elemental map

For **ET4** she did not reconsider the endmembers found in VisuaLayered and indicated that she was very confident in the identified ones. In the end, she did not go over the last task, **ET5**, in either software as establishing the order of paint layers is a more complex process that requires more time than we had available and additional, different data (e.g. the SWIR data cube). The only observation that she made during the ENVI based analysis was that the pigment lead white appears to be applied over ochre. She observed this when creating the endmember maps and annotated the corresponding map to indicate this. Since **ET5** also refers to the annotation of pixels according to the pigment they represent (**T2.4**), even without the indication of paint layer order, our collaborator saved the identified pigment names as part of the ENVI and VisuaLayered endmember map names.

6.2. Survey Results

In this section we briefly discuss the answers provided in the survey by our collaborator. Her profession is conservation scientist (Q1) and she is familiar with both data modalities and their combined analysis. The questions about the five analysis tasks and about the comparison between the two workflows were answered on a 5-point Likert scale [76] (Figure 6.11). In the end, we asked 5 general feedback questions, followed by the 10 Standard Usability questions [77].

Section	Statement	Score 1	2	3	4	5
Tasks 1,2	2. The image view is efficient for exploring areas of interest in the painting					●
	3. Dragging the RGB lines in the spectrum view and seeing real-time updates of the image RGB channels helped me identify hidden structures in the painting					
	4. The embedding view complemented the image view and helped me identify pixels with similar spectra					●
	5. The standard deviation area around an average spectrum improved the selection of endmembers *					●
	6. I can easily customize and filter the saved endmembers in the spectrum view					●
	7. The linking of the image, embedding and spectrum views improved the endmember identification process					●
	8. Updating the threshold angle for SAM in the spectrum view and seeing instant changes in the endmember map helped in deciding the best threshold angle for a map *			●		
Tasks 3,4	9. Linking the two datasets and using the image and embedding views for each data modality helps me understand the relation between the two data modalities					●
	10. The embedding of the MA-XRF data helped me better understand the elemental composition of similar pixels *			●		
	11. The cluster correlation view helped me identify pixels that are similar with respect to both data modalities					●
	12. The cluster correlation view improved my understanding of the molecular and elemental composition of pixels *					●
	13. Using the embedding of one data modality with the color overlay as the clustering of the other modality helps me find similarities between the two data modalities					●
	14. Using the embedding of one data modality with the color overlay of one dimension from the other data modality allows me to focus on a detail of one data modality and analyze it in the context of the other one					●
	15. It was useful to explore the pixels with a known elemental composition by looking at two dimensions of the MA-XRF dataset in the embedding view *			●		
	16. I could easily establish the elemental composition of endmembers using the cluster correlation view				●	
	17. The combination of different MA-XRF and RIS maps in the image view is not useful for understanding the spatial overlap of different endmembers and chemical elements					
Task 5	18. Using the cluster correlation view with the endmember and chemical element clusters helps me better understand the order of the paint layers					
	19. Combining the two data modalities in one embedding view does not help in understanding the order of the paint layers					
Comparison	21. Compared to ENVI, I could more efficiently **statement** in VisuaLayered:					
	explore the RIS data	●				
	explore the MA-XRF data					●
	identify endmembers				●	
	explore the spatial distribution of endmembers		●			
	explore the spatial distribution of pixels with similar elemental composition					●
	jointly explore pixels with similar spectral and elemental profiles					●
	reason about the order of paint layers					
	classify pixels according to the pigment they represent				●	
General Feedback	22. I am confident in the pigments that I found by following the tasks in the evaluation					●

Figure 6.11: Collaborator answers to survey questions on a 5-point Likert scale from *strongly disagree* (1) to *strongly agree* (5); The statements followed by * are rephrased to their positive form in this table and their scores are inverted for representation purposes

For the first two tasks (Q2-Q8), she was overall happy with the workflow for endmember identification (e.g. the use of the embedding view for identifying similar spectra or the linking of the image, embedding and spectrum views for finding endmembers). For Q8, she indicated that the setting of the threshold angle for SAM could be improved. Moreover, she did not provide any feedback on the interactive setting of the RGB wavelengths (Q3) as she did not use it (they already looked at false color images in ENVI). For the next two tasks (Q9-Q17), they were again happy with the provided functionalities for pigment identification (e.g. the use of the cluster correlation view for identifying pixels that are similar with respect to both data modalities or the linking of the two datasets across all views). Since she did not combine any maps in the image view, Q17 was left unanswered. Since the last task was not completed, there are no questions answered for this task (Q18-Q20).

When comparing the two workflows (Q21), she indicated that she could more efficiently explore the RIS data and the spatial distribution of endmembers using ENVI. However, she considered VisuaLayered better for identifying endmembers. Moreover, for the exploration of the MA-XRF data and the combined analysis, she considered VisuaLayered more efficient. Lastly, she considered VisuaLayered better for identifying pigments.

For the general feedback (Q22-Q26), she wrote that VisuaLayered helps her understand the correlation of the two data modalities, ”in particular with mixtures”. She indicated that she is very confident in the identified pigments. For the best features of the system, she mentioned: the embedding of the RIS data for easier endmember identification, the embedding of one modality with the color overlay of a cluster (set) for ”navigating the dataset” and the cluster correlation view with the endmember and the elemental maps for having ”a final identification of the pigments”. Moreover, she wrote what she thinks

is missing from the system: a histogram for deciding the threshold angle for SAM (similar to the one in ENVI), the visualization of the spectra in the MA-XRF data cube (before it is transformed into elemental maps) and the computation of the first derivative of the RIS data cube. Furthermore, she mentioned that it was difficult to adjust the minimum threshold using the slider in the cluster per dimension settings. Lastly, the SUS score is 85 out of 100.

6.3. Discussion

In Section 6.3.1 we discuss additional results obtained based on the evaluation and we go over the limitations of the evaluation in Section 6.3.2.

6.3.1. Additional Results

At the end of the evaluation, we had a short discussion with our collaborator about other ways of using the views as part of the analysis. We considered the embedding view with the color overlay of all the endmember map clusters (Figure 6.7a). Our collaborator used the color overlay of only one endmember at various moments in the analysis, like in Figure 6.7b. However, by using all the endmember map clusters, we can explore the distribution of all endmembers in the MA-XRF embedding and use this view for selecting one particular endmember, that we want to further explore.

Additionally, we discussed about visualizing two dimensions of the MA-XRF data, so the combination of two chemical elements, in the embedding view (Figure 6.12). This could be used for identifying pixels that represent a pigment which is characterized by two chemical elements. Moreover, we could see a clear separation between these pixels and pixels which contain only one of the two elements. Hence, this view can be used for **T2.1** if only the MA-XRF data is available and for **T2.3** if we also use the endmember map clusters as the color overlay, as shown in Figure 6.12a. In this case, the selected points correspond to *end2*, so we can see the points that contain only Cu and the points that contain Cu and Co, just like in Figure 6.7.

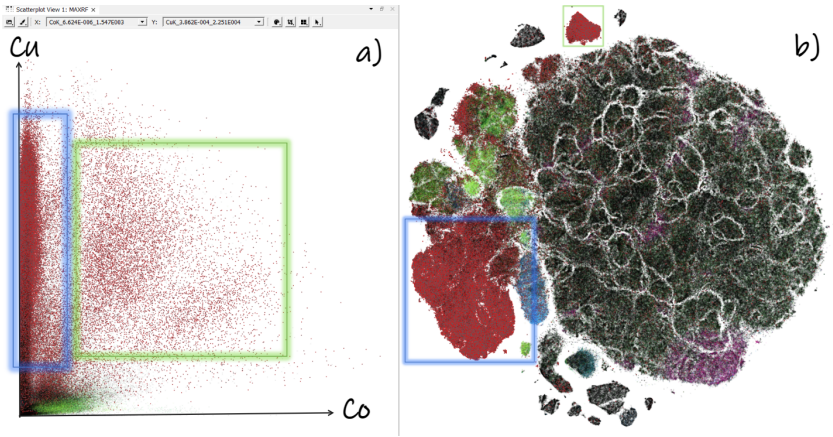

Figure 6.12: Overlap between *end2* map cluster and two selected chemical elements;

Additionally, after the evaluation we observed that the vermilion map (Figure 6.13a) obtained in ENVI is better than the map obtained in VisuaLayered (Figure 6.13b). We know this, because similarly to the leaf in Figure 6.10, the flower highlighted in Figure 6.13c contains mercury, which is present in vermilion and hence the flower should be highlighted in its endmember map. Upon further analysis, we observed that the flower appears as a separate cluster in the embedding (indicated in red in Figure 6.14a) and its average spectrum is a lighter version of the vermilion endmember chosen in VisuaLayered (shown in blue in Figure 6.14c). If we combine both the endmember and the flower spectrum SAM maps, as shown in Figure 6.14d, we obtain a very similar map to the one inFigure 6.13a. Hence, we can still obtain the same result in VisuaLayered and have a distinction between the lighter and darker vermilion. Alternatively, we can discard the initially selected endmember and save just the flower spectrum as an endmember. We can see it is very similar to the ENVI endmember (shown in purple in Figure 6.13c). If we slightly increase the threshold angle for SAM, we also obtain an almost identical map to the ENVI one.

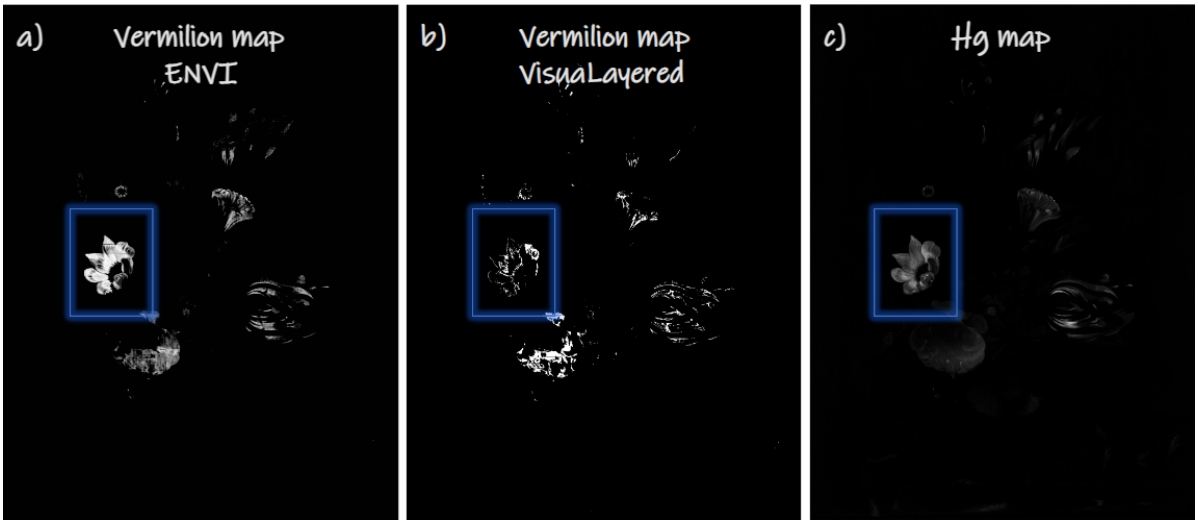

Figure 6.13: Vermilion map comparison; a) ENVI endmember map; b) VisuaLayered endmember map; c) Hg elemental map

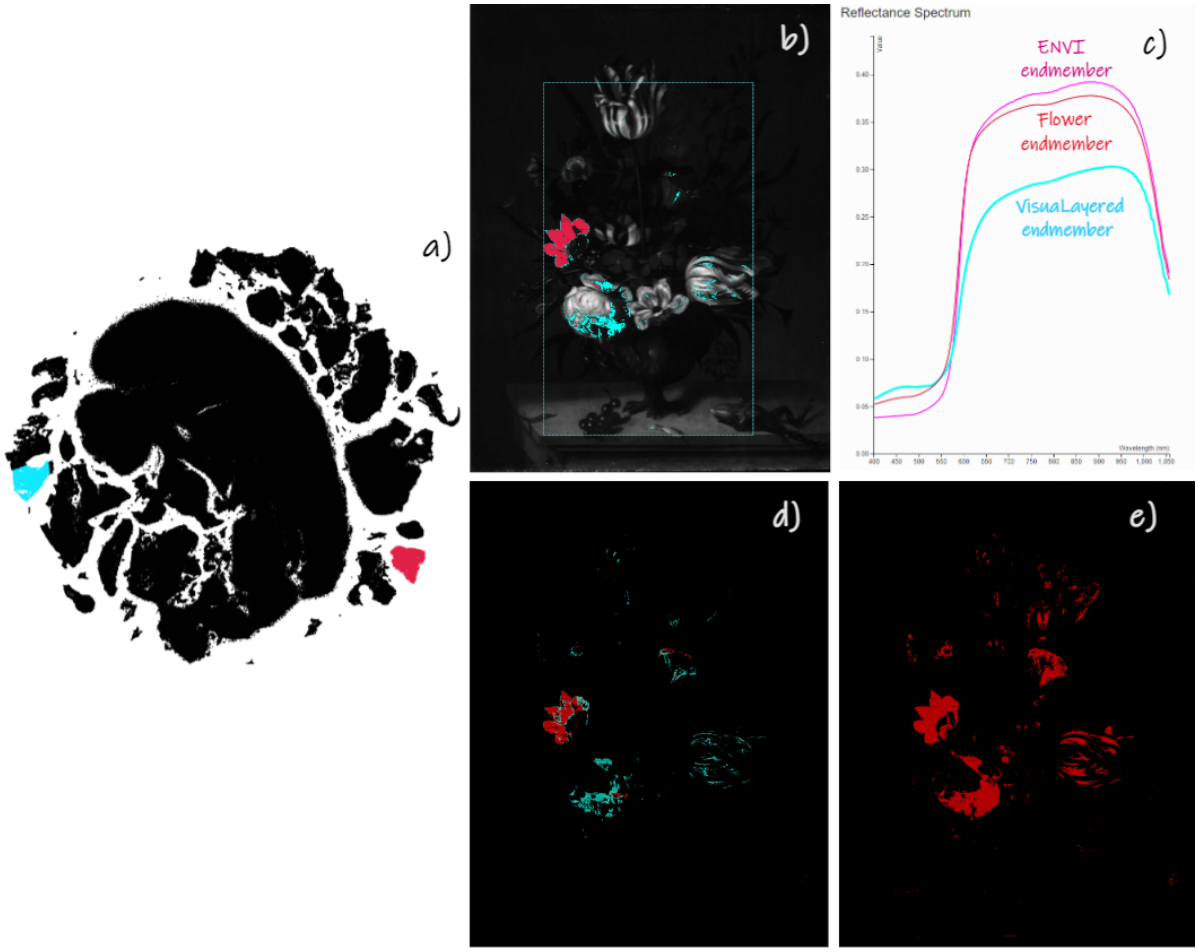

Figure 6.14: A better mapping of vermilion; a) RIS embedding showing the endmember cluster (in blue) and flower cluster (in red); b) The same two clusters highlighted in the Image View; c) The two cluster spectra and the ENVI endmember in the Spectrum View; d) The VisuaLayered endmember and flower spectrum maps combined in the Image View; e) The flower endmember map

Our collaborator used some of the automatically selected endmembers for mapping in ENVI, while in VisuaLayered she only used the automatic clustering to get an initial idea of the spectra present in the image. In the end, however, most of the maps for the automatically found endmembers were

discarded in ENVI. This indicates that the manual selection, based on domain expert knowledge, is the most relevant one, indicating that our initial reasons to provide an interactive workflow are valid. Additionally, the 6 maps obtained in ENVI and VisuaLayered for endmembers that represent the same pigment are very similar (e.g. the azurite maps in Figure 6.10a,b). Therefore, we wanted to also compare the corresponding sets of endmembers and we loaded them in a spectrum view in VisuaLayered (Figure 6.15). We set the same hue and saturation for every two endmembers representing one pigment, with the VisuaLayered endmembers in a lower lightness. We can observe that the VisuaLayered endmembers have similar shapes to the ones in ENVI, with peaks at the same wavelengths, but at different values on the y-axis. In general, the VisuaLayered endmembers are darker than the ENVI ones, with the exception of azurite.

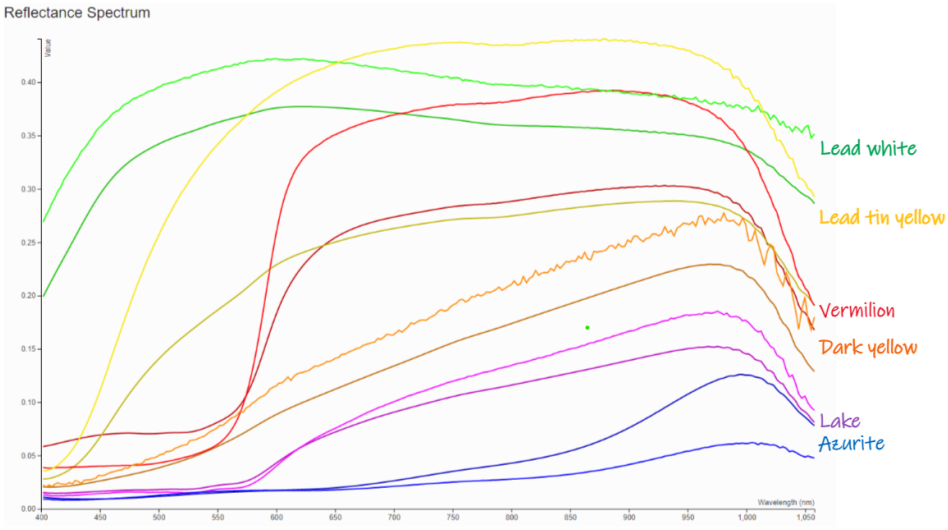

Figure 6.15: Comparison of endmembers representing the same pigment; a) ENVI endmembers; b) VisuaLayered endmembers

Lastly, due to the fact that t-SNE uses the cosine distance in order to project the high-dimensional data into 2D, which is the same distance used in the Spectral Angle Mapper, the clustering of the embedding could be used to replace the SAM computation for all the endmembers. Using SAM, we map points that have a similar spectrum to a selected endmember. In the t-SNE embedding, points that have a similar spectrum are grouped together. We create clusters based on these groups of points and save their average spectra as endmembers. Ideally, the saved endmembers have a small standard deviation. Hence, the points in an endmember cluster are points with a similar spectrum to an endmember, just like the points in an endmember map obtained with SAM.

We explored this idea using the evaluation data. As a starting point in our example, we use the automatic clustering of the t-SNE embedding . We over-clustered the embedding into 133 clusters and manually refined them by either discarding small, separate clusters (of less than 10000 points), merging clusters that were part of one single group of points and editing clusters that had an average spectrum with a large standard deviations. After refinement, we were left with 35 clusters (Figure 6.16d), which mostly have average spectra with a low standard deviation. We use the cluster correlation view to compare these clusters with the 19 endmember map clusters obtained by our collaborator. We discuss two examples based on this clustering.

The first example is shown in Figure 6.16. We picked and further analysed an endmember map (of *end19*) that mostly overlaps with only one of the clusters (*cluster 10*), as shown in Figure 6.16d. The points from the *end19* map cluster are selected and shown in red in the embedding, mostly as part of the pink cluster, which is *cluster 10*. The map of *end19* shows the pigment present in the table. The endmember and the cluster spectrum (Figure 6.16a) are very similar, but the cluster contains more pixels representing the table, than the *end19* map created by our collaborator (the map and the cluster points are highlighted in red in Figure 6.16c). As the standard deviation of *cluster 10* is not larger than that of *end19*, the cluster could be a better representation of the endmember. Additionally, the points of the *end19* map that do not overlap with *cluster 10*, belong to its neighbouring cluster (also highlighted in Figure 6.16d). This cluster contains points with a darker version of the endmember. Hence, we could

also consider a combined map of the two clusters instead of the endmember map.

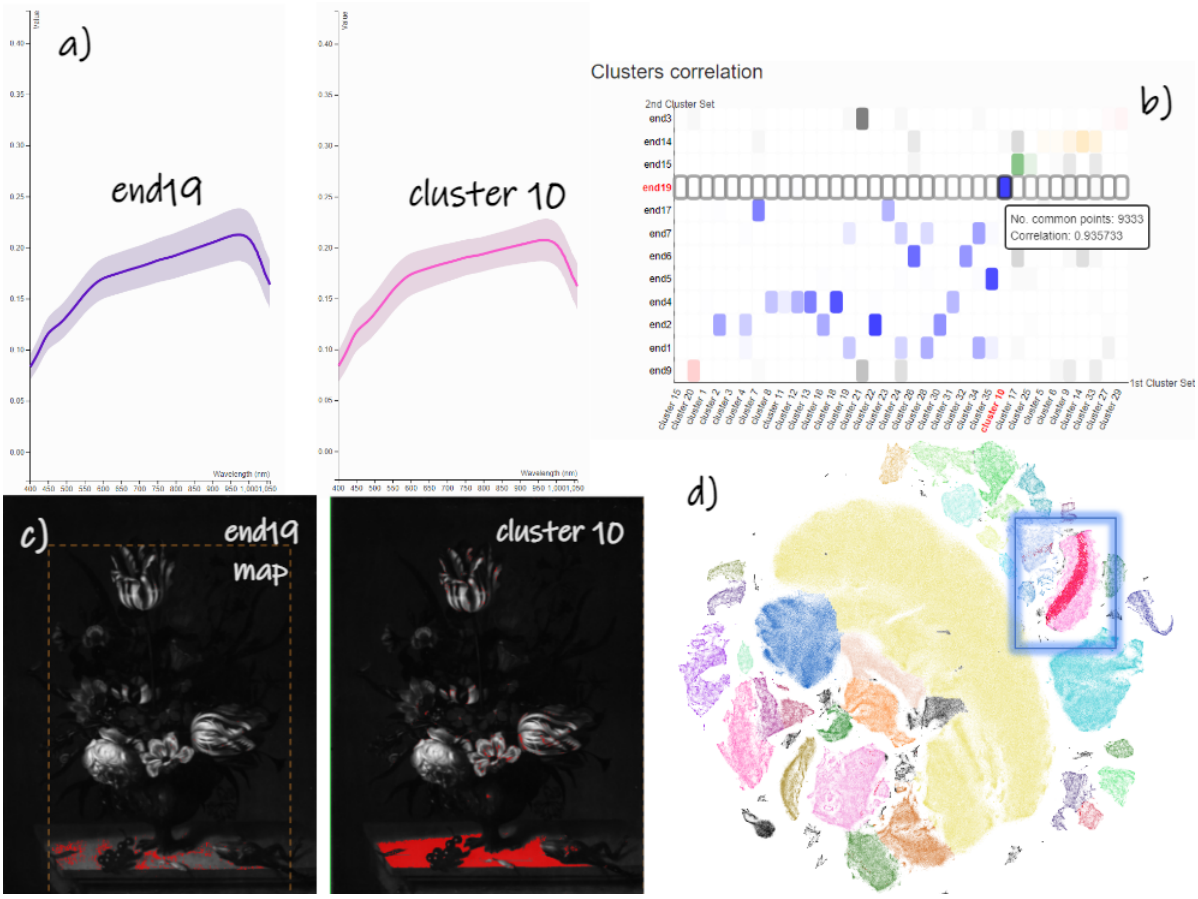

Figure 6.16: Example 1; a) *end19* and *cluster 10* spectrum; b) Correlation of automatically found clusters (x-axis) and endmember map clusters (y-axis) with *end19* map cluster selected; c) Points in endmember map and cluster highlighted in two image views; d) RIS embedding with color overlay of automatically found clusters and *end19* map points selected and shown in *cluster 10* (pink cluster)

The second example is shown in Figure 6.17. For this example we picked the *end2* map, that overlaps with multiple clusters. In the embedding we selected the *end2* map points. We discuss *clusters 22 and 30* as they have the biggest overlap with the map. The spectra of the two clusters and the endmember are similar, with different values for the peaks (Figure 6.17). Since the *end2* map contains most of the points in *cluster 22* and *end2* has a smaller standard deviation than the *cluster 22* spectrum, we assume that the other clusters that are selected, like *cluster 30* (that has a darker spectrum), contribute to the reduction of the standard deviation. Having the different clusters provides us with more information on the different spectra representing one pigment, as opposed to creating just one endmember map. We can again consider a combined map of the clusters (Figure 6.17c) and refine the clusters based on the elemental maps in case too many/few points are selected (similarly to choosing the threshold angle for SAM). The clustering of the t-SNE embedding does not require setting a threshold value and, if done automatically, it attributes one cluster to every point in the dataset. It can be a more direct way of obtaining maps of the endmembers. However, domain knowledge and the MA-XRF elemental maps are still required in order to refine and obtain the most representative clusters for the existing pigments.

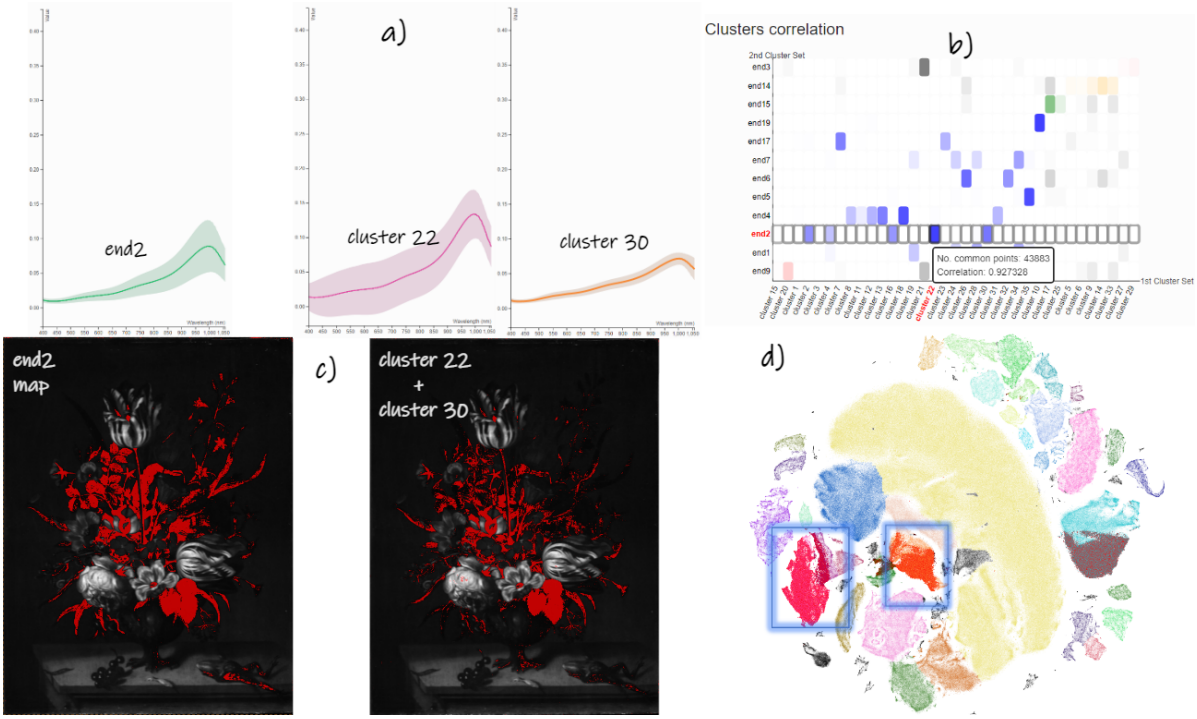

Figure 6.17: Example 2; a) *end2*, *cluster 22* and *cluster 30* spectra; b) Correlation of automatically found clusters (x-axis) and endmember map clusters (y-axis) with *end2* map cluster selected; c) Points in endmember map and cluster highlighted in two image views; d) RIS embedding with color overlay of automatically found clusters and *end2* map points selected and shown in *cluster 22* (pink cluster) and *cluster 30* (orange cluster)

6.3.2. Limitations

We start by mentioning that the evaluation was done with only one participant, which is not enough for proving the efficiency and usability of VisuaLayered. This case study is a pilot and a refined version of the five analysis tasks and survey can be used for a larger study. It was difficult to find other domain experts with similar levels of experience with both data modalities, who could test our proposed workflow in the available time. Hence, we view the case study described in Section 6.1 as an example use of VisuaLayered that shows the functionalities of the system in relation to the proposed workflow. Moreover, we recognize that the evaluation is not completely impartial as our collaborator took part in the development process of VisuaLayered and provided feedback on possible improvements along the way. However, we took this into consideration when creating the survey and alternated affirmative and negative statements in the survey. During the evaluation, we also mostly observed her work and only answered questions when something was unclear, without making any suggestions on what visualizations should be used.

We decided to start the comparison in the evaluation with the analysis in ENVI and afterwards in VisuaLayered, so that we can observe whether VisuaLayered improves our collaborator's typical workflow. Moreover, we did not want her to use the results obtained in VisuaLayered as a guide for her analysis in ENVI. However, this still occasionally happened the other way around. For example, she had a previously identified endmember in mind and tried to also find it using VisuaLayered. However, we think the comparison still produced valuable results and as mentioned in Section 6.1, she even found a few additional endmembers using VisuaLayered.

The t-SNE computation requires some processing time. For this reason we decided to compute the embeddings of the MA-XRF and RIS datasets before the evaluation. It took us around three hours to run t-SNE on both datasets on the hardware specified in Figure 6.18. However, this is not necessarily problematic, as we can perform the t-SNE computation once per dataset and then use the embeddings throughout the entire analysis.

Additionally, in an ideal setting the MA-XRF and RIS datasets would be completely aligned. We characterize the pigments in the painting with respect to both data modalities by looking at pixel values in both data cubes. Hence, corresponding pixels in the two datasets should represent the same spatial position in the painting. However, the scanning is done with two devices with different specifications

Category	Component
OS	Windows 10
CPU	AMD Ryzen 7 2700x (8 Cores, 16 Logical Processors, 3700MHz Clock Frequency)
GPU	NVIDIA GeForce RTX 2070 SUPER (8GB GDDR6 VRAM)
RAM	32GB DDR4, 3200 MHz

Figure 6.18: Hardware specifications used for running t-SNE

and spectral sampling distances and the resulting data cubes are not completely aligned. Moreover, the MA-XRF dataset is obtained by stitching together multiple smaller data cubes of parts of the painting, so the alignment within the dataset is also not perfect. Furthermore, we can observe that noise is present in the data (for example in the spectrum of one single pixel) due to the real life conditions of scanning. By considering average spectra of the RIS data and by refining the chemical element clusters of the MA-XRF data, we also aim to reduce the noise in the data.

7

Future Work

For future work, we consider possible additions to the analysis workflow based on the evaluation results and related work.

We would like to automatically set a basis threshold angle for the SAM computation in order to assist the user. Our collaborator uses a histogram in ENVI which shows the number of similar pixels according to a given threshold value. They normally set the threshold around the first change in the values in the histogram. We could identify this change automatically, set the threshold right before the change and generate the map. Then, the user could still interactively change the threshold in a similar histogram view. Furthermore, during the evaluation we learned that determining whether a spectrum is an endmember, depends on the location of its peaks. Hence, we would like to add the possibility of zooming in on user selected wavelength ranges of the spectrum in the spectrum view. Moreover, users should then be able to use the selected range for the SAM computation, such that they can identify pixels that are similar to an endmember only based on the significant part of the endmember. Additionally, based on the qualitative feedback received in the evaluation, we want to make minor usability improvements such as decreasing the slider step sizes for the SAM/SCM threshold and the cluster per dimension minimum/maximum thresholds, reducing the computation time of the SAM/SCM maps and improving the interaction in the histogram view.

Next, we consider techniques that could be added to the workflow based on related work. The joint data cube of the two data modalities' raw cubes might offer a new perspective on the available data. However, combining the two datasets is not straightforward. Firstly, the two datasets would have to be completely aligned, which is also a current requirement of the datasets analysed in VisuaLayered. Secondly, the datasets represent different high-dimensional spaces and simply fusing the RIS and MA-XRF spectra for every scanned pixel would lead to a loss of information. Thirdly, neither of the two datasets should dominate the other one in computations on the fused cube. Alfeld et al. [3] fused the RIS-XRF and MA-XRF cubes for their application. They aligned the RIS dataset, that has a higher lateral resolution, to the MA-XRF dataset and chose relevant spectral ranges in both datasets to balance their influence on the fused cube. They obtained a cube with 784 channels, from which 59% of the channels represent the XRF data and 41% represent the RIS data. Using the t-SNE embedding of the combined cube, they identify pigments and present their results. We would like to include a fused data cube in the current workflow of VisuaLayered and use it to complement the combined analysis of the MA-XRF and RIS datasets. By computing the t-SNE embedding of the combined cube, we would have an additional space for clustering the pixels and we could explore the correlation of these clusters to the clusters of the two data modalities in order to better understand the fused space.

Furthermore, spectral unmixing methods for RIS data could provide more insight into the composition of mixed pigments. We could use automatic methods, like Maximum Distance (MaxD) [78] or Simplex Identification via Split Augmented Lagrangian (SISAL) [8] in order to help domain experts identify the endmembers present in the painting. This would be similar to using the Spectral Hourglass Wizard in ENVI for automatically identifying endmembers. Additionally, we would consider automatic pigment identification methods, which work by matching the scanned reflectance spectra to a library of pigment endmembers in order to identify pigments, possibly without the need to first identify endmembers in the painting. However, since the reflectance spectrum of a pigment is influenced by the surface

and ground layer of a painting, we would like to use an endmember library that is created based on similar works to the investigated painting [78]. In VisuaLayered, we could add such automatic methods as an initial step in the analysis, so that we identify a set of endmembers, which experts can validate and manually refine. Moreover, we could use methods that provide a confidence value on the identified endmembers, such that experts could start refining the endmembers which have low confidence values.

Lastly, we would like to explore the visualization of a time series of RIS datasets [31]. The pigments used in a painting degrade over time because of numerous reasons, like external factors (e.g. light, temperature). Hence, if we perform multiple RIS scans of one painting, over the years, we might be able to see how the pigments change by comparing the data from different scans. By using a variant of t-SNE, Joint t-SNE [79], we obtain comparable projections across time steps. Joint t-SNE keeps the points that do not change over time in the same location across projections and hence clearly shows the points that change and move in the projections. Therefore, it could help us understand what points represent pigments that changed their molecular structure over time. We could use the embedding view and its functionalities to show the projections and interact with the data.

We believe a structured evaluation with a greater number of domain experts is the next step for assessing the efficiency of VisuaLayered and obtaining quantitative results as well. As an introduction for the evaluation, we could prepare a demo showing an example analysis workflow in VisuaLayered and a document describing all the functionalities of the system, such that participants can learn how to use the system even if we are not assisting them. Furthermore, in order to make the evaluation shorter with respect to time and therefore accessible to more participants, we could provide a simplified survey and a list of simpler tasks that participants would have to complete using only VisuaLayered. For example, for the first task instead of asking participants to identify all endmembers in a painting, we could ask them to identify a fixed number endmembers. Moreover, we could set a time limit per task and take into account how far every participant gets in the available time. Moreover, we should consider different datasets obtained with or based on RIS that are normally used as part of the analysis, such as SWIR data cubes and the first derivative of RIS data cubes. We would like to add the computation of the first derivative data cube as an additional plugin. Moreover, our collaborator suggested the addition of the MA-XRF original data cube (before its processing into elemental maps) in the visual analysis. If we change the axes domain in the spectrum view, we could also visualize the MA-XRF spectra and analyse them in combination with the other data.

8

Conclusion

We designed and implemented VisuaLayered, an interactive workflow for the combined visual analysis of Reflectance Imaging Spectroscopy and Macro X-Ray Fluorescence data for pigment identification in paintings. VisuaLayered is an integrated system which consists of five views that allow the user to explore and interact with these two complementary data modalities. This work was done in collaboration with the Rijksmuseum in Amsterdam. We conducted a field study with our collaborator and learned about their typical analysis workflow in order to design our workflow. In VisuaLayered the user can explore the spatial distribution of pixels with different elemental and molecular compositions in the image view, identify and compare endmembers by using the embedding and spectrum views and better understand the correlation between the two data modalities using the cluster correlation and embedding views. The linking of all available views offers a direct connection between the two data modalities throughout the entire analysis process. VisuaLayered is meant to help domain experts in their analysis and hence, interactivity is an important aspect of all the available views.

We tested the functionalities of VisuaLayered in a case study. We asked our collaborator to complete five painting analysis tasks using both their typical analysis tools and using VisuaLayered. The aim of the case study was to obtain an example analysis workflow in VisuaLayered and discuss how well it matches the proposed workflow and whether it improves our collaborator's typical workflow. We learned that the visual combination of the two data modalities represents an improvement as it allows them to better and more easily identify the connections between the MA-XRF and RIS data. In particular, they found the embedding view with different color overlays and the cluster correlation view of endmember maps and chemical element clusters really useful. Additionally, clustering the embedding with the help of the spectrum view, which shows the standard deviation of selected average spectra, improved the process of endmember identification. In the end, they found more endmembers using VisuaLayered.

Reflection To conclude, my thesis was a very interesting, fun and at times challenging project. I think making a planning at the beginning of the thesis really helped me organize my thoughts about the project into explicit actions that I could to take in order to complete my project. The planning also helped me understand the extent of my project. At the beginning, I could have planned more meetings with our collaborators in order to have an even clearer picture of the cultural heritage background. During the development of the prototype, I could have found one or two more domain experts who were willing to test the prototype so I would have more feedback on what can be improved. I did not always meet the deadlines I set for myself in the planning, but even with small delays I was generally able to complete all the planned tasks. However, I did not work on the report consistently throughout the project (as indicated in the planing), so this led to me having a harder time when it came to writing the report. Lastly, if I prepared the evaluation earlier in the process, we could have asked a few more people to participate in the evaluation.

A

Appendix

A.1. Short Paper

The following paper was submitted and accepted as a short paper at the Eurographics Workshop on Graphics and Cultural Heritage 2022. This paper is based on my thesis work and focuses on the visual analysis of RIS data for endmember identification and analysis. I would like to thank all the co-authors of the paper for their feedback and contribution.

EUROGRAPHICS Workshop on Graphics and Cultural Heritage (2022)
R. Pintus and F. Ponchio (Editors)

Visual Analysis of RIS Data for Endmember Selection

A. Popa[1], F. Gabrieli[2], T. Kroes[3], A. Krekeler[2], M. Alfeld[1], B. Lelieveldt[3], E. Eisemann[1], and T. Höllt[1]

[1]Delft University of Technology, Delft, The Netherlands
[2]Conservation and Science Department, Rijksmuseum, Amsterdam, The Netherlands
[3]Leiden University Medical Center, Leiden, The Netherlands

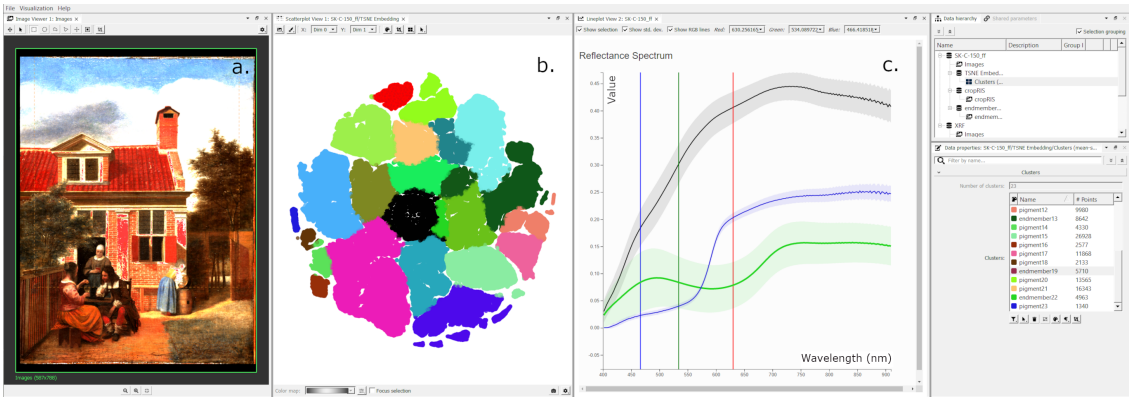

Figure 1: ***Overview of the implemented system*** *showing the painting* Figures in a Courtyard behind a House *by Pieter De Hooch [Kre19]; The pixel selection is highlighted in red in the a) image view, b) embedding view and shown as the top spectrum in the c) spectrum view.*

Abstract

Reflectance Imaging Spectroscopy (RIS) is a hyperspectral imaging technique used for investigating the molecular composition of materials. It can help identify pigments used in a painting, which are relevant information for art conservation and history. For every scanned pixel, a reflectance spectrum is obtained and domain experts look for pure representative spectra, called endmembers, which could indicate the presence of particular pigments. However, the identification of endmembers can be a lengthy process, which requires domain experts to manually select pixels and visually inspect multiple spectra in order to find accurate endmembers that belong to the historical context of an investigated painting. We propose an integrated interactive visual-analysis workflow, that combines dimensionality reduction and linked visualizations to identify and inspect endmembers. Here, we present initial results, obtained in collaboration with domain experts.

CCS Concepts

• ***Human-centered computing*** → ***Visual analytics; Visualization systems and tools;*** • ***Applied computing*** → *Fine arts;*

1. Introduction

Digital, non-invasive imaging techniques provide new insights into cultural heritage. Domain experts register and analyze properties of materials present in paintings, in order to identify pigments and guide the conservation and restoration of old master paintings. Moreover, this information is relevant for art history to enable dating and establishing the authenticity of a painting [BPP05]. One such technique is Reflectance Imaging Spectroscopy (RIS), a hyperspectral technique which captures the molecular composition of the material being scanned as reflectance spectra. Experts look for pure representative spectra, called *endmembers* [RDF*12], which correspond to different pigments, allowing them to establish the pigment distribution in a painting. Domain experts can manually select pixels and analyze their spectra in order to find pigments, but this can be a lengthy process. Automatic endmember selection [GMGM18; MBG*22] shows promising results for pigment identification, but depends on matching spectra to a set of predefined endmembers. Domain experts often need to identify wrongly classified pixels and to help improve the automatic endmember selection process. Further, there is limited comparability between endmembers derived from different paintings, as other factors such as aging of the painting, measurement batches, etc. have an impact on the acquired data. Therefore, domain expert knowledge is crucial in the analysis process.

We propose a visual analysis system that keeps domain experts in the analysis loop and provides semi-automatic methods for endmember detection. This work is the first part of an ongoing design study, conducted in collaboration with domain experts in order to create a visual analysis workflow for painting analysis. In particular, here, we focus on the interactive extraction of endmembers. The main contributions of this work are:

- a visual analysis workflow for interactive identification of endmembers in RIS data
- a prototype implementation of the workflow

While endmember selection is an essential part of the analysis process, more than just molecular information is needed to characterize a pigment or a combination of pigments and understand the layered composition of a painting. In the broader scope of this project, we aim to integrate the analysis of RIS data with other imaging modalities, such as Macro X-Ray Fluorescence (MA-XRF) which is commonly used to inspect the elemental distribution in paint layers. [DDVV20; MTS*19]

2. Related Work

Visual exploration of high-dimensional data is a vast area in visualization research. An in-depth review is out of the scope of this paper. A good overview is presented in the survey by Liu et al. [LMW*17], while Brehmer et al. [BSIM14] present an overview of typical applications, domains, and corresponding tasks.

Recently, high-dimensional imaging data has become more widely available in several application domains, including cultural heritage, and consequently a number of visual analysis approaches are available. Grabowski et al. [GMGM18] use t-SNE [vdMH] as a base for clustering pixels and then automatically identify pigments by matching spectra to predefined endmembers representing different pigments. Pouyet et al. [PRK*18] compare t-SNE to PCA and MNF for visualizing hyperspectral data of paint samples and obtain the best results for pigment identification with t-SNE. Alfeld et al. [APMW18] applied t-SNE on RIS and MA-XRF data. With the help of the t-SNE embedding visualized in a scatterplot, they were able to identify most of the pigments present in the wall paintings of an Egyptian tomb. We build on these developments as one aspect of our proposed workflow.

Hyper3D [KRF*14] is a visualization system designed to aid art conservators, focused on hyperspectral volumetric data. It provides a pixel-wise reflectance spectrum view for detail exploration of the data. We follow a similar approach of developing an integrated system, but focus on paint layer information revealed by RIS where the spectrum view offers a broader range of analysis options for endmember identification.

3. Requirements Analysis

In order to establish the requirements for the system, we conducted a series of meetings with collaborating imaging-science experts from the Rijksmuseum, Amsterdam. After discussing their general data analysis needs, we conducted a field study consisting of semi-structured interviews and observed their current workflow for analyzing RIS data. After implementing our initial prototype, we have deployed the software with our partners and have conducted regular follow-up meetings to identify issues and provide updates.

https://www.overleaf.com/project/62c96f03f5a459ae819940dd

Data Reflectance Imaging Spectroscopy (RIS) is a hyperspectral imaging technique that goes beyond standard photography by measuring continuous spectra of reflected light for each imaged pixel [APMW18]. Depending on the acquisition hardware, different ranges of the light spectrum are captured. For example, VNIR (visible to near-infrared spectroscopy) cameras capture the range of 400 nm to 1000 nm and SWIR (short-wave infrared spectroscopy) cameras capture 1000 nm to 2500 nm. The data is stored as a volume where the x- and y-dimensions correspond to the spatial extents of the data and the z-dimension to a discrete sampling of the spectral information. In the following, we interpret the spectral information as a high-dimensional space and thus the data at each x, y position as a high-dimensional pixel.

Throughout this paper, we use a VNIR dataset (covering roughly 400 nm to 900 nm) of the painting *Figures in a Courtyard Behind a House* by Pieter de Hooch [Kre19]. The data cube consists of 1174×1756 pixels with a pixel size of roughly 0.35×0.35 mm^2. Each pixel represents a sampling of the spectral information with 2.54 nm, resulting in 200 dimensions.

Current Workflow The current endmember extraction workflow of our collaborators is based on a combination of tools and scripts around the geospatial analysis software ENVI [L3H]. They start by selecting different areas in the painting in an image viewer and then inspect the corresponding spectra in a lineplot. Once they have an initial overview of the data, they manually and/or automatically define endmembers using the Spectral Hourglass Wizard (ENVI-SHW). Lastly, they create endmember maps by plotting points with similar spectra to an endmember using the Spectral Angle Mapper [DM00]. At this point they can look at the maps and reason about the used pigments.

Task abstraction Based on the data and observed workflow, we have identified the following tasks that the proposed system must support. The user

T1: *identifies* different endmember candidates in the high-dimensional image data,
T2: *compares* and *filters* identified endmember candidates based on their spectral information, and
T3: *explores* the spatial distribution of pixels relating to identified endmembers and mapping parameters.

4. Proposed Solution

Based on the requirements analysis (Section 3), we created the design and implemented a prototype of our proposed system that we refined during follow-up meetings with our collaborators.

4.1. Design

The proposed system is shown in Figure 1. It consists of three main views to support the tasks described in Section 3. All views in the system are linked to enable comparison across all views.

4.1.1. Image View

The first view (Figure 1a) is an image view. The view can be used in different modes to support tasks **T1** and **T3**. For probing the image for potential endmembers (**T1**), the view is used to show the original image data. Since showing the full spectral information is not

feasible on an RGB-screen, we provide different modes; if a registered photograph of the painting is available, this can be shown directly. Additionally, the user can select any wavelength to show the corresponding scalar field, optionally as false colors using a colormap, or select three different wavelengths and map them to the red, green, and blue channels of the screen. Figure 1a shows an example using three wavelengths in the visible red, green, and blue spectra, inidcated by the vertical lines in Figure 1c, emulating a photograph. However, channels can easily be remapped, for example to show information of the otherwise invisible infrared part of the spectrum. The user can interactively make selections of pixels in this view and visualize the corresponding spectra in the spectrum view (Section 4.1.3).

Additionally, the image view is also used for **T3**. Any created endmember map (Section 4.1.3) can be shown in the view, either individually or as a combination, e.g., by mapping up to three different maps to the RGB channels (Figure 2).

4.1.2. Embedding View

Figure 1b shows the embedding view. We use t-SNE [vdMH] to create an embedding of the original VNIR data. To support the calculation of t-SNE embeddings on more than two million pixels we use a GPU-based implementation of t-SNE [PTM*20]. Every point in the embedding corresponds to a pixel in the image, where points with a similar spectrum are placed close to each other. We use the cosine distance to calculate the pairwise distances between the discrete spectra. The t-SNE view is used to identify groups of similar pixels that can serve as potential endmembers (**T1**), without relying on the visual inspection in image space. The user can probe the embedding via a selection and visualize the corresponding spectra in the spectrum view (Section 4.1.3). This task can also be supported by clustering the embedding using Mean Shift clustering [HPvU*16] and using the resulting clusters to derive endmembers.

4.1.3. Spectrum View

The third view (Figure 1c) is the spectrum view that is implemented as a lineplot, where the x-axis represents the wavelengths in *nm* and the y-axis represents the measured reflectance. The view shows the mean spectrum of any selection of pixels from the two other views as a black line. In addition to the mean, it is possible to visualize the standard deviation of a spectrum as an area around the line (e.g., colored areas in Figure 2). When the user has identified a suitable spectrum for an endmember, the selection can be made persistent in this view (**T2**). Persistent spectra are added to a list view and can optionally be visualized in the lineplot with a user-defined color.

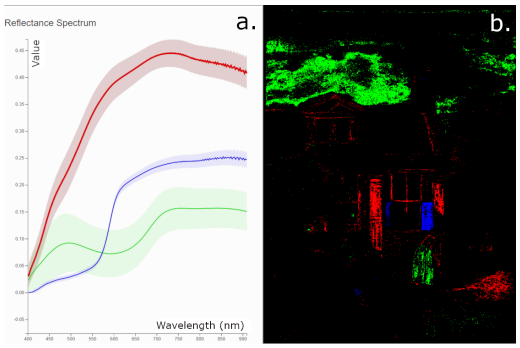

Figure 2: ***Endmember maps*** *corresponding to three endmembers.*

To allow further inspection of the image, based on the spectra, the user can set the wavelengths for coloring the image view directly in this view. Three vertical lines (Figure 1c) represent the R, G, and B channels of the image view and can be dragged along the x-axis to any combination of interest. For example, the user can set the values of the three lines in the near-infrared range to identify areas that are perceptually similar, but differ in their spectral composition. The change is reflected immediately in the image view so that the user can gradually see how their wavelength selections influence the colors in the image.

As described above, when a selection is made persistent to create an endmember, it is added to a list view on the side of the spectrum view. From this view, the user can toggle the creation of an endmember map for each item, using the Spectral Angle Mapper (SAM) (**T3**). SAM computes the cosine distance of the endmember to all pixels in the image. We choose SAM based on the current workflow of our collaborators, but optionally provide mapping using the Spectral Correlation Mapper [DM00; DGH14]. For either, the distance is typically thresholded to create binary maps, where similar pixels are set to true and all others to false. Figure 2 shows an image combining three endmember maps, obtained with SAM, using the R, G, and B channels.

4.2. Implementation

We implemented our workflow in a plugin-based framework for high-dimensional data analysis. The system is implemented in Qt/C++ and visualizations in OpenGL and D3 [BOH11], respectively, according to the required performance. Upon completion of the full project, we plan to release the tool as open source.

5. Preliminary Results

We have deployed an initial prototype of the implemented system with our domain expert collaborators. Based on their feedback, we have gone through several iterations to add functionality.

Our collaborators have successfully identified endmembers using the system in initial testing. In their early feedback, they were enthusiastic about the integrated analysis workflow. The linked views with real-time selections make the analysis easier and offer a better understanding of the data.

In this first phase, the focus was on interactive, manual exploration of the data. Initial experiments lead us to believe that clustering based on the t-SNE embedding can further improve the process and potentially also replace the SAM computation. We use the same cosine distance metric as SAM when computing the pairwise distances in t-SNE. As a result, obtained clusters contain points that are similar to an endmember in an SAM map. Figure 3 shows a visual comparison of an endmember map created using SAM and one based directly on a cluster extracted from the t-SNE map. The points belonging to the endmember cluster (Figure 3b) are highlighted in red and the points selected through SAM for the same endmember (Figure 3c) are highlighted in blue. Based on a threshold angle, SAM selected dissimilar points, which are scattered in other clusters than the endmember cluster (Figure 3d) and which

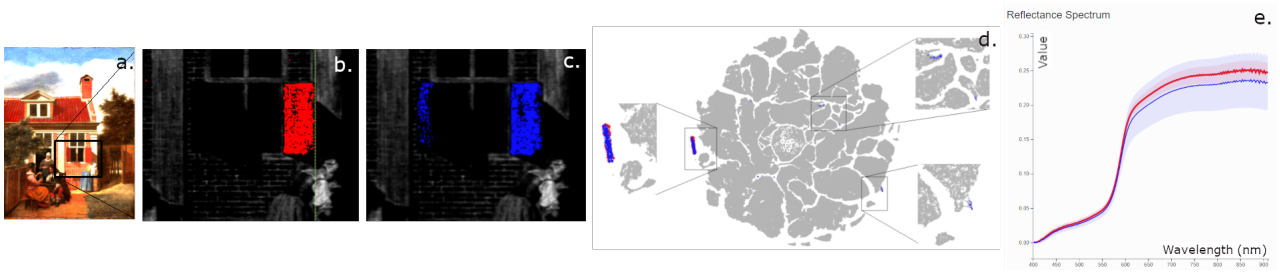

Figure 3: ***Comparison*** *of t-SNE cluster (red) and endmember-based (blue) map.*

lead to a larger standard deviation of the SAM points visible in blue in the spectrum view (Figure 3e). t-SNE offers a more automatic solution than SAM as it does not require predefined endmembers or setting threshold angles for identifying similar spectra. We plan a structured comparison in future work.

6. Conclusion and Future Work

We presented an integrated, interactive system and analysis workflow for endmember identification and mapping. Initial results and feedback from our collaborators show that the system improves their current workflow. As part of our research, we want to further develop the system to allow for the combined visual analysis of MA-XRF and RIS data. In particular, we aim to provide means to identify points with similar elemental composition and study correlations between the two data modalities. Further, it would be interesting to explore ways to model pixels as a combination of endmembers using spectral unmixing [GTG21] as pigments can appear in a painting as such mixtures.

References

[APMW18] ALFELD, MATTHIAS, PEDETTI, SILVIA, MARTINEZ, PHILIPPE, and WALTER, PHILIPPE. "Joint data treatment for Vis–NIR reflectance imaging spectroscopy and XRF imaging acquired in the Theban Necropolis in Egypt by data fusion and t-SNE". *Comptes Rendus Physique* 19.7 (2018), 625–635 2.

[BOH11] BOSTOCK, MICHAEL, OGIEVETSKY, VADIM, and HEER, JEFFREY. "D^3 data-driven documents". *IEEE transactions on visualization and computer graphics* 17.12 (2011), 2301–2309 3.

[BPP05] BARNI, MARIO, PELAGOTTI, ANNA, and PIVA, ALESSANDRO. "Image processing for the analysis and conservation of paintings: opportunities and challenges". *IEEE Signal processing magazine* 22.5 (2005), 141–144 1.

[BSIM14] BREHMER, MATTHEW, SEDLMAIR, MICHAEL, INGRAM, STEPHEN, and MUNZNER, TAMARA. "Visualizing dimensionally-reduced data: Interviews with analysts and a characterization of task sequences". *Proceedings of ACM BELIV Workshop*. 2014, 1–8 2.

[DDVV20] DELANEY, JOHN K, DOOLEY, KATHRYN A, VAN LOON, ANNELIES, and VANDIVERE, ABBIE. "Mapping the pigment distribution of Vermeer's Girl with a Pearl Earring". *Heritage Science* 8.1 (2020), 1–16 2.

[DGH14] DEBORAH, HILDA, GEORGE, SONY, and HARDEBERG, JON YNGVE. "Pigment mapping of the scream (1893) based on hyperspectral imaging". *International Conference on image and Signal processing*. Springer. 2014, 247–256 3.

[DM00] DE CARVALHO, O ABILIO and MENESES, PAULO ROBERTO. "Spectral correlation mapper (SCM): an improvement on the spectral angle mapper (SAM)". *Summaries of the 9th JPL Airborne Earth Science Workshop, JPL Publication 00-18*. Vol. 9. 2000 2, 3.

[GMGM18] GRABOWSKI, BARTOSZ, MASARCZYK, WOJCIECH, GŁOMB, PRZEMYSŁAW, and MENDYS, AGATA. "Automatic pigment identification from hyperspectral data". *Journal of Cultural Heritage* 31 (2018), 1–12. ISSN: 1296-2074. DOI: https://doi.org/10.1016/j.culher.2018.01.003. URL: https://www.sciencedirect.com/science/article/pii/S1296207417306544 1, 2.

[GTG21] GRILLINI, FEDERICO, THOMAS, JEAN-BAPTISTE, and GEORGE, SONY. "Comparison of Imaging Models for Spectral Unmixing in Oil Painting". *Sensors* 21.7 (2021), 2471 4.

[HPvU*16] HÖLLT, THOMAS, PEZZOTTI, NICOLA, van UNEN, VINCENT, et al. "Cytosplore: Interactive immune cell phenotyping for large single-cell datasets". *Computer Graphics Forum*. Vol. 35. 3. 2016, 171–180 3.

[Kre19] KREKELER, A. "A Study of Pieter de Hooch's Painting Technique". *Pieter de Hooch in Delft. From the Shadow of Vermeer*. W-Books, 2019, 56–79 1, 2.

[KRF*14] KIM, MIN H, RUSHMEIER, HOLLY, FFRENCH, JOHN, et al. "Hyper3D: 3D graphics software for examining cultural artifacts". *Journal on Computing and Cultural Heritage (JOCCH)* 7.3 (2014), 1–19 2.

[L3H] L3HARRIS TECHNOLOGIES, INC. *L3Harris ENVI*. https://www.l3harris.com/all-capabilities/envi. Accessed: 2022-06-23 2.

[LMW*17] LIU, SHUSEN, MALJOVEC, DAN, WANG, BEI, et al. "Visualizing High-Dimensional Data: Advances in the Past Decade". *IEEE Transactions on Visualization and Computer Graphics* 23.3 (2017), 1249–1268 2.

[MBG*22] MELADA, JACOPO, BONIZZONI, LETIZIA, GARGANO, MARCO, et al. "Pre-processing methods for automatic pigment recognition through Vis–NIR reflectance spectra". *The European Physical Journal Plus* 137.1 (2022), 44 1.

[MTS*19] MACLENNAN, DOUGLAS, TRENTELMAN, KAREN, SZAFRAN, YVONNE, et al. "Rembrandt's An Old Man in Military Costume: Combining hyperspectral and MA-XRF imaging to understand how two paintings were painted on a single panel". *Journal of the American Institute for Conservation* 58.1-2 (2019), 54–68 2.

[PRK*18] POUYET, EMELINE, ROHANI, NEDA, KATSAGGELOS, AGGELOS K, et al. "Innovative data reduction and visualization strategy for hyperspectral imaging datasets using t-SNE approach". *Pure and Applied Chemistry* 90.3 (2018), 493–506 2.

[PTM*20] PEZZOTTI, NICOLA, THIJSSEN, JULIAN, MORDVINTSEV, ALEXANDER, et al. "GPGPU Linear Complexity t-SNE Optimization". *IEEE Transactions on Visualization and Computer Graphics (Proceedings of IEEE VAST 2019)* 1 (2020), 1172–1181 3.

[RDF*12] RICCIARDI, PAOLA, DELANEY, JOHN K, FACINI, MICHELLE, et al. "Near infrared reflectance imaging spectroscopy to map paint binders in situ on illuminated manuscripts". *Angewandte Chemie International Edition* 51.23 (2012), 5607–5610 1.

[vdMH] Van der MAATEN, LAURENS and HINTON, GEOFFREY. "Visualizing data using t-SNE". *Journal of machine learning research* 9.11 (), 2579–2605 2, 3.

A.2. Evaluation Survey

Combined Visual Analysis of MA-XRF and RIS Data - Evaluation

You are being invited to participate in a research study titled Combined Visual Analysis of MA-XRF and RIS Data. This study is done by Andra Popa from the TU Delft.

The purpose of this research study is to evaluate the workflow within our proposed system, VisuaLayered, in comparison to another existing system, the geospatial analysis software ENVI.

We prepared a list of painting analysis tasks for you to follow. You start by completing the tasks in ENVI, then in VisuaLayered. For our software, the tasks are followed by a few questions. Once you finish a task, please fill in its corresponding questions. In the end, we prepared a few questions regarding the comparison of the two systems and a few general feedback questions regarding our software.

In case any of the questions are not applicable to your workflow, please leave them unanswered.

The painting used to show example views in our software is *Figures in a Courtyard behind a House* by Pieter de Hooch (A. Krekeler, catalogue of the exhibition Pieter de Hooch in Delft. From the Shadow of Vermeer, A. Jansen (ed.), W-Books, Zwolle, 2019, pp. 56-79).

The form has 8 sections. The questions will take approximately 20 minutes to complete. The data will be used for assessment of our system and reported in a thesis report and defence.

As with any online activity the risk of a breach is always possible. To the best of our ability your answers in this study will remain confidential. We will minimize any risks by not recording your name.

Your participation in this study is entirely voluntary and you can withdraw at any time. You are free to omit any questions. The research data will be stored 10 years or more, in accordance with the TU Delft Research Data Framework Policy.

Corresponding Researcher:
Andra Popa
a.popa-2@student.tudelft.nl

Personal Information

1

What is your profession?

Evaluation Tasks - ENVI

Please complete the following tasks using **ENVI**:

1. Identify the endmembers present in the studied painting.

2. Explore the spatial distribution of pixels that are spectrally similar to the found endmembers.

3. Consider the endmembers from task 1 (and the pixels that are similar to them) and establish their elemental composition as well, in order to identify the used pigments.

4. Are you satisfied with your endmembers? If not, consider refining your endmember selections.

5. Can you reason about the possible order of the paint layers? If yes, annotate your selections to indicate this (name clusters e.g. *top_ultramarine*, *bottom_ochre*).

Evaluation Tasks - VisuaLayered

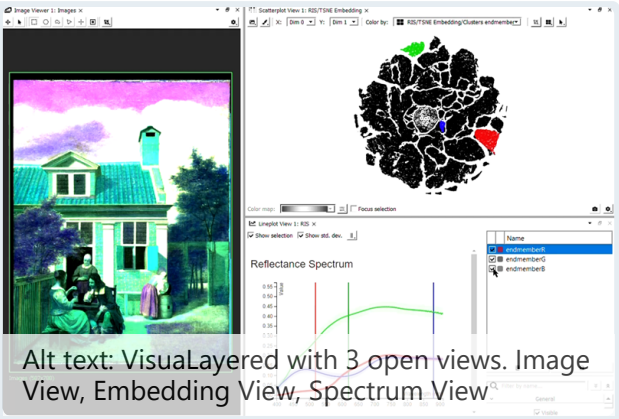

Alt text: VisuaLayered with 3 open views. Image View, Embedding View, Spectrum View

Please complete the following tasks using **VisuaLayered**:

1. Identify the endmembers present in the studied painting.

2. Explore the spatial distribution of pixels that are spectrally similar to the found endmembers.

Example workflow using the RIS dataset

1. You can identify areas of interest in the painting using the image view and explore the average spectra of your selections in the spectrum view (Figures 1, 2). The image display range can be set using window/level settings. You can interactively set these settings in the histogram view by dragging the ends of the diagonal green line (once the window/level fields from the two views are linked to each other) to match the data range (Figure 3). Moreover, you can interactively set the RGB channels of the image by dragging the three RGB lines in the spectrum view, once the channels and the lines are linked to each other (Figure 6c).

2. Next, you can explore the spatial distribution of spectrally similar pixels in the t-SNE embedding of the data in the embedding view (Figure 4a). Points with similar spectra are grouped together. You can use the spectrum view to observe the average spectra of your selections and look for potential endmembers with a low standard deviation area around a spectrum (Figure 4b). You can save a selection as a cluster in the embedding view. Alternatively, you can automatically cluster the entire dataset using Mean Shift Clustering and manually refine this clustering. You can color the points in the embedding based on the clustering (Figure 5).

3. Then, you can load a set of clusters in the spectrum view as endmembers (Figure 6c). The spectrum view allows you to change the color and name of an endmember (Figure 7a).

4. For every found endmember, you can perform the Spectral Angle Mapper (SAM) algorithm in the spectrum view in order to obtain an endmember map (Figure 7b). You can interactively set the threshold angle for SAM in this view (Figure 7a).

5. You can explore a combination of endmember maps or endmember clusters in the image view (Figure 8). A combination of 3 maps can be obtained by setting one map per RGB channel of the image.

Once you completed the tasks, please answer the following questions. The questions should take around 2 minutes to answer.

2

The image view is efficient for exploring areas of interest in the painting.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

3

Dragging the RGB lines in the spectrum view and seeing real-time updates of the image RGB channels helped me identify hidden structures in the painting.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

4

The embedding view complemented the image view and helped me identify pixels with similar spectra.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

5

The standard deviation area around an average spectrum did not improve the selection of endmembers.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

6

I can easily customize and filter the saved endmembers in the spectrum view.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

7

The linking of the image, embedding and spectrum views improved the endmember identification process.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

8

Updating the threshold angle for SAM in the spectrum view and seeing instant changes in the endmember map did not help in deciding the best threshold angle for a map.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

Evaluation Tasks - VisuaLayered

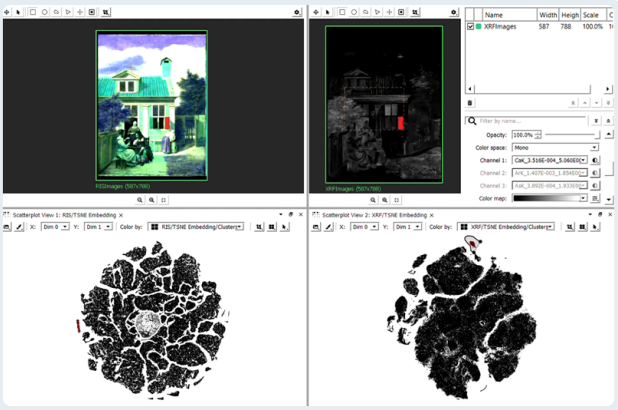

Please complete the following tasks in **VisuaLayered**:

3. Consider the endmembers from task 1 (and the pixels that are similar to them) and establish their elemental composition as well, in order to identify the used pigments.

4. Are you satisfied with your endmembers? If not, consider refining your endmember selections.

Example workflow using the MA-XRF and RIS datasets

1. You can start by grouping together the two datasets. This ensures the linking of all open views that show data from either dataset. You can use the image and embedding views like in the previous example, but this time also for the MA-XRF data in order to observe the similarities between the two datasets (Figure 9). You can then also cluster the MA-XRF data to obtain groups of points that have a similar elemental composition.

2. Next, you can observe the correlation of both cluster sets in the cluster correlation view (Figure 10a).

3. Moreover, you can combine the two data modalities in one embedding view. You can have the embedding of one data modality open and color the points based on the clustering of the other modality (Figure 11) or color the points based on the distribution of one dimension from the other modality (Figure 12).

4. If you want to look at two specific chemical elements' distribution in combination, you can also load the MA-XRF dataset in the embedding view and choose the two dimensions/elements you want to visualize (Figures 13, 14).

5. You can also save the RIS and MA-XRF maps as clusters (by choosing the *Cluster per Dimension* option for each dataset and manually changing the obtained clusters) and also load them in the cluster correlation view in order to see the elemental composition of an endmember (Figure 15). You can also cluster all the points that do not belong to any cluster in a cluster set (by choosing the *Complement cluster* option) in order to potentially identify areas in the painting that are unexplored (Figure 16).

6. In the end, you can explore the overlap between any combination of maps from both data modalities. Any number of maps can be combined by loading, in the image view, one map per dataset and setting the opacity level to 50% and a constant color per map (Figures 17, 18).

Once you completed the tasks, please answer the following questions. The questions should take around 2 minutes to answer.

9

Linking the two datasets and using the image and embedding views for each data modality helps me understand the relation between the two data modalities

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

10

The embedding of the MA-XRF data did not help me better understand the elemental composition of similar pixels.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

11

The cluster correlation view helped me identify pixels that are similar with respect to both data modalities.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

12

The cluster correlation view did not improve my understanding of the molecular and elemental composition of pixels.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

13

Using the embedding of one data modality with the color overlay as the clustering of the other modality helps me find similarities between the two data modalities.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

14

Using the embedding of one data modality with the color overlay of one dimension from the other data modality allows me to focus on a detail of one data modality and analyze it in the context of the other one.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

15

It was not useful to explore the pixels with a known elemental composition by looking at two dimensions of the MA-XRF dataset in the embedding view.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

16

I could easily establish the elemental composition of endmembers using the cluster correlation view.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

17

The combination of different MA-XRF and RIS maps in the image view is not useful for understanding the spatial overlap of different endmembers and chemical elements.

	1	2	3	4	5	
Strongly agree	○	○	○	○	○	Strongly disagree

Evaluation Tasks - VisuaLayered

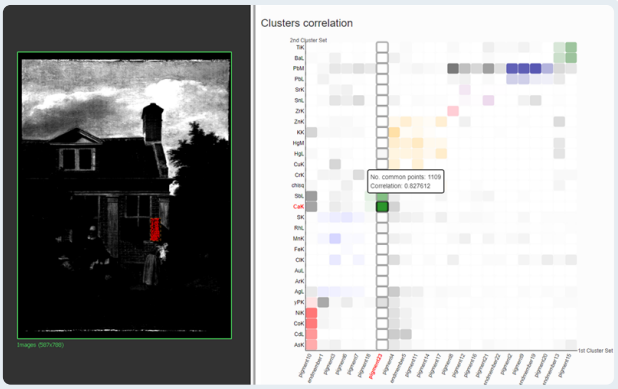

Please complete the following tasks in **VisuaLayered**:

5. Can you reason about the possible order of the paint layers? If yes, annotate your selections to indicate the order (name clusters e.g. top_ultramarine, bottom_ochre).

For this task, you can use the combined views described in the previous section.

Once you completed the task, please answer the following questions. The questions should take around 2 minutes to answer.

18

Using the cluster correlation view with the endmember and chemical element clusters helps me better understand the order of the paint layers.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

19

Combining the two data modalities in one embedding view does not help in understanding the order of the paint layers.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

20

How did you use the available views in order to complete this task?

Comparison between ENVI and VisuaLayered

Now that you completed the tasks in both systems, we would like to ask you to compare them.

This part should take around 3 minutes to complete.

21

Please choose a score for the following statements:

Compared to ENVI, I could more efficiently **statement** **in VisuaLayered:**

	ENVI is more efficient		same efficiency		**VisuaLayered** is more efficient
explore the RIS data	○	○	○	○	○
explore the MA-XRF data	○	○	○	○	○
identify endmembers	○	○	○	○	○
explore the spatial distribution of endmembers	○	○	○	○	○
explore the spatial distribution of pixels with similar elemental composition	○	○	○	○	○
jointly explore pixels with similar spectral and elemental profiles	○ **ENVI** is more efficient	○	○ same efficiency	○	○ **VisuaLayered** is more efficient

reason about the order of paint layers	◯	◯	◯	◯	◯
classify pixels according to the pigment they represent	◯	◯	◯	◯	◯

General Feedback

This section contains 5 questions about our system. This part should take around 5 minutes to complete.

22

I am confident in the pigments that I found by following the tasks in the evaluation.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

23

Do you use the MA-XRF data in order to identify endmembers? If the answer is yes, please explain why you use it.

24

Does the system help you understand the relation between the two data modalities? Please explain your answer.

25

What is the best feature of the system?

26

What is missing from the system and would help you in your analysis?

General Usability of the Software

Please answer the following questions regarding our system.

This part should take around 2 minutes to complete.

27

I think that I would like to use this system frequently.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

28

I found the system unnecessarily complex.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

29

I thought the system was easy to use.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

30

I think that I would need the support of a technical person to be able to use this system.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

31

I found the various functions in this system were well integrated.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

32

I thought there was too much inconsistency in this system.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

33

I would imagine that most people would learn to use this system very quickly.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

34

I found the system very cumbersome to use.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

35

I felt very confident using the system.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

36

I needed to learn a lot of things before I could get going with this system.

	1	2	3	4	5	
Strongly disagree	○	○	○	○	○	Strongly agree

Bibliography

[1] M. Barni, A. Pelagotti, and A. Piva, "Image processing for the analysis and conservation of paintings: Opportunities and challenges," *IEEE Signal Processing Magazine*, vol. 22, no. 5, pp. 141–144, 2005. DOI: 10.1109/MSP.2005.1511835 (cit. on p. 1).

[2] A. Martins *et al.*, "Jazz colors: Pigment identification in the gouaches used by henri matisse," *Heritage*, vol. 4, no. 4, pp. 4205–4221, 2021. DOI: 10.3390/heritage4040231 (cit. on p. 1).

[3] M. Alfeld, S. Pedetti, P. Martinez, and P. Walter, "Joint data treatment for vis–nir reflectance imaging spectroscopy and xrf imaging acquired in the theban necropolis in egypt by data fusion and t-sne," *Comptes Rendus Physique*, vol. 19, no. 7, pp. 625–635, 2018 (cit. on pp. 1, 3, 8, 49).

[4] L. van der Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. 11, pp. 2579–2605, 2008 (cit. on pp. 1, 8, 21).

[5] S. Liu, D. Maljovec, B. Wang, P.-T. Bremer, and V. Pascucci, "Visualizing high-dimensional data: Advances in the past decade," *IEEE transactions on visualization and computer graphics*, vol. 23, no. 3, pp. 1249–1268, 2016. DOI: 10.1109/TVCG.2016.2640960 (cit. on p. 1).

[6] S. A. Barcellos Lins, B. Bremmers, and G. E. Gigante, "Xismus — x-ray fluorescence imaging software for multiple samples," *SoftwareX*, vol. 12, p. 100 621, 2020. DOI: 10.1016/j.softx.2020.100621 (cit. on pp. 1, 7).

[7] M. Picollo, A. Casini, C. Cucci, J. Jussila, M. Poggesi, and L. Stefani, "Chapter a new compact vnir hyperspectral imaging system for non-invasive analysis in the fineart and architecture fields," 2018. DOI: 10.36253/978-88-6453-707-8.16 (cit. on pp. 1, 7).

[8] B. Grabowski, W. Masarczyk, P. Głomb, and A. Mendys, "Automatic pigment identification from hyperspectral data," *Journal of Cultural Heritage*, vol. 31, pp. 1–12, 2018. DOI: 10.1016/j.culher.2018.01.003 (cit. on pp. 1, 5, 8, 49).

[9] M. Sedlmair, M. Meyer, and T. Munzner, "Design study methodology: Reflections from the trenches and the stacks," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 12, pp. 2431–2440, 2012. DOI: 10.1109/TVCG.2012.213 (cit. on p. 2).

[10] P. Shippert, "Introduction to hyperspectral image analysis," *Online Journal of Space Communication*, vol. 2, no. 3, p. 8, 2003 (cit. on p. 3).

[11] S. Prasad, J. Chanussot, J. E. Fowler, J. Bioucas-Dias, and C. D. Creuserre, "Introduction to the issue on advances in hyperspectral data processing and analysis," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 6, pp. 961–963, 2015. DOI: 10.1109/JSTSP.2015.2457631 (cit. on pp. 3, 8).

[12] M. L. Clarke *et al.*, "Imaging spectroscopies to characterize a 13th century japanese handscroll, the miraculous interventions of jizō bosatsu," *Heritage Science*, vol. 9, no. 1, pp. 1–10, 2021. DOI: 10.1186/s40494-021-00497-1 (cit. on p. 3).

[13] N. Keshava and J. F. Mustard, "Spectral unmixing," *IEEE signal processing magazine*, vol. 19, no. 1, pp. 44–57, 2002. DOI: 10.1109/79.974727 (cit. on p. 4).

[14] M. Alfeld and K. Janssens, "Strategies for processing mega-pixel x-ray fluorescence hyperspectral data: A case study on a version of caravaggio's painting supper at emmaus," *Journal of analytical atomic spectrometry*, vol. 30, no. 3, pp. 777–789, 2015. DOI: doi.org/10.1039/C4JA00387J (cit. on pp. 4, 12).

[15] M. S. Shackley, "An introduction to x-ray fluorescence (xrf) analysis in archaeology," in *X-ray fluorescence spectrometry (XRF) in geoarchaeology*, Springer, 2011, pp. 7–44. DOI: 10.1007/978-1-4419-6886-9_2 (cit. on p. 4).

[16] A. Locquet, J. Dong, M. Melis, and D. Citrin, “Terahertz pulsed imaging reveals the stratigraphy of a seventeenth-century oil painting,” in *Unconventional Optical Imaging*, vol. 10677, 2018, pp. 373–376. DOI: 10.1117/12.2306771 (cit. on p. 4).

[17] J. Dik, K. Janssens, G. Van Der Snickt, L. Van Der Loeff, K. Rickers, and M. Cotte, “Visualization of a lost painting by vincent van gogh using synchrotron radiation based x-ray fluorescence elemental mapping,” *Analytical chemistry*, vol. 80, no. 16, pp. 6436–6442, 2008. DOI: 10.1021/ac800965g (cit. on p. 4).

[18] J. Orsilli, A. Galli, L. Bonizzoni, and M. Caccia, “More than xrf mapping: Steam (statistically tailored elemental angle mapper) a pioneering analysis protocol for pigment studies,” *Applied Sciences*, vol. 11, no. 4, p. 1446, 2021. DOI: 10.3390/app11041446 (cit. on p. 5).

[19] A. Galli and L. Bonizzoni, “Contribution of x-ray fluorescence techniques in cultural heritage materials characterization,” *Applied Sciences*, vol. 12, no. 13, p. 6309, 2022. DOI: 10.3390/books978-3-0365-4868-5 (cit. on p. 5).

[20] A. Mazzinghi *et al.*, “Ma-xrf for the characterisation of the painting materials and technique of the entombment of christ by rogier van der weyden,” *Applied Sciences*, vol. 11, no. 13, p. 6151, 2021. DOI: 10.3390/app11136151 (cit. on p. 5).

[21] M. Picollo, A. Casini, C. Cucci, L. Stefani, R. Jiménez-Garnica, and L. Fuster-López, “Documentation and analysis of some picasso’s paintings by using hyperspectral imaging technique to support their conservation and stylistic matters,” in *IOP Conference Series: Materials Science and Engineering*, vol. 949, 2020, p. 012 023. DOI: 10.1088/1757-899X/949/1/012023 (cit. on p. 5).

[22] M. Martínez Domingo, E. Valero, R. Huertas, M. Durbán, T. Espejo, and R. Blanc, “Spectral information to get beyond color in the analysis of water-soluble varnish degradation,” *Heritage Science*, vol. 7, no. 1, pp. 1–16, 2019. DOI: 10.1186/s40494-019-0323-5 (cit. on p. 5).

[23] N. Richard, D. Helbert, C. Olivier, and M. Tamisier, “Pseudo-divergence and bidimensional histogram of spectral differences for hyperspectral image processing,” *Journal of Imaging Science and Technology*, vol. 60, no. 5, pp. 50 402–1, 2016. DOI: 10.2352/J.ImagingSci.Technol.2016.60.5.050402 (cit. on p. 5).

[24] F. Grillini, J.-B. Thomas, and S. George, “Comparison of imaging models for spectral unmixing in oil painting,” *Sensors*, vol. 21, no. 7, p. 2471, 2021. DOI: 10.3390/s21072471 (cit. on pp. 5, 7).

[25] E. D’Elia *et al.*, “Macro x-ray fluorescence and vnir hyperspectral imaging in the investigation of two panels by marco d’oggiono,” *Microchemical Journal*, vol. 154, p. 104 541, 2020. DOI: 10.1016/j.microc.2019.104541 (cit. on pp. 5, 8).

[26] M. Alfeld, M. Mulliez, J. Devogelaere, L. de Viguerie, P. Jockey, and P. Walter, “Ma-xrf and hyperspectral reflectance imaging for visualizing traces of antique polychromy on the frieze of the siphnian treasury,” *Microchemical Journal*, vol. 141, pp. 395–403, 2018. DOI: 10.1016/j.microc.2018.05.050 (cit. on p. 6).

[27] J. K. Delaney, K. A. Dooley, A. Van Loon, and A. Vandivere, “Mapping the pigment distribution of vermeer’s girl with a pearl earring,” *Heritage Science*, vol. 8, no. 1, pp. 1–16, 2020. DOI: 10.1186/s40494-019-0348-9 (cit. on p. 6).

[28] V. A. Solé, E. Papillon, M. Cotte, P. Walter, and J. Susini, “A multiplatform code for the analysis of energy-dispersive x-ray fluorescence spectra,” *Spectrochimica Acta Part B: Atomic Spectroscopy*, vol. 62, no. 1, pp. 63–68, 2007. DOI: 10.1016/j.sab.2006.12.002 (cit. on p. 7).

[29] P. Campos *et al.*, “A low-cost portable system for elemental mapping by xrf aiming in situ analyses,” *Applied Radiation and Isotopes*, vol. 152, pp. 78–85, 2019. DOI: 10.1016/j.apradiso.2019.06.018 (cit. on p. 7).

[30] M. Cotte, T. Fabris, G. Agostini, D. Motta Meira, L. De Viguerie, and V. A. Solé, “Watching kinetic studies as chemical maps using open-source software,” *Analytical chemistry*, vol. 88, no. 12, pp. 6154–6160, 2016. DOI: 10.1021/acs.analchem.5b04819 (cit. on p. 7).

[31] D. Langenkämper *et al.*, "Exploring time series of hyperspectral images for cold water coral stress response analysis," *PloS one*, vol. 17, no. 8, e0272408, 2022. DOI: 10.1371/journal.pone.0272408 (cit. on pp. 7, 8, 50).

[32] M. H. Kim, H. Rushmeier, J. Ffrench, I. Passeri, and D. Tidmarsh, "Hyper3d: 3d graphics software for examining cultural artifacts," *Journal on Computing and Cultural Heritage (JOCCH)*, vol. 7, no. 3, pp. 1–19, 2014. DOI: 10.1145/2567652 (cit. on p. 7).

[33] L3Harris Technologies, Inc., *L3harris envi*, https://www.l3harris.com/all-capabilities/envi, Accessed: 2022-06-23 (cit. on pp. 8, 13).

[34] S. Liu, D. Maljovec, B. Wang, P.-T. Bremer, and V. Pascucci, "Visualizing high-dimensional data: Advances in the past decade," *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 3, pp. 1249–1268, 2017. DOI: 10.1109/TVCG.2016.2640960 (cit. on p. 8).

[35] S. WOLD, K. ESBENSEN, and P. GELADI, "Principal component analysis," *Chemometrics and Intelligent Laboratory Systems*, vol. 2, pp. 37–52, 1987. DOI: 10.1016/0169-7439(87)80084-9 (cit. on p. 8).

[36] A. J. Izenman, "Linear discriminant analysis," in *Modern Multivariate Statistical Techniques: Regression, Classification, and Manifold Learning*. 2008, pp. 237–280. DOI: 10.1007/978-0-387-78189-1_8 (cit. on p. 8).

[37] L. McInnes, J. Healy, and J. Melville, "Umap: Uniform manifold approximation and projection for dimension reduction," *arXiv preprint arXiv:1802.03426*, 2018. DOI: 10.48550/arXiv.1802.03426 (cit. on p. 8).

[38] W. S. Torgerson, "Multidimensional scaling: I. theory and method," *Psychometrika*, vol. 17, no. 4, pp. 401–419, 1952 (cit. on p. 8).

[39] E.-a. D. Amir *et al.*, "Visne enables visualization of high dimensional single-cell data and reveals phenotypic heterogeneity of leukemia," *Nature biotechnology*, vol. 31, no. 6, pp. 545–552, 2013. DOI: 10.1038/nbt.2594 (cit. on p. 8).

[40] W. Li, J. E. Cerise, Y. Yang, and H. Han, "Application of t-sne to human genetic data," *Journal of bioinformatics and computational biology*, vol. 15, no. 04, p. 1 750 017, 2017. DOI: 10.1142/S0219720017500172 (cit. on p. 8).

[41] A. Vieth, A. Vilanova, B. Lelieveldt, E. Eisemann, and T. Höllt, "Incorporating texture information into dimensionality reduction for high-dimensional images," in *2022 IEEE 15th Pacific Visualization Symposium (PacificVis)*, IEEE, 2022, pp. 11–20. DOI: 10.48550/arXiv.2202.09179 (cit. on p. 8).

[42] B. Melit Devassy and S. George, "Dimensionality reduction and visualisation of hyperspectral ink data using t-sne," *Forensic Science International*, vol. 311, p. 110 194, 2020. DOI: 10.1016/j.forsciint.2020.110194 (cit. on p. 8).

[43] B. Melit Devassy, S. George, and P. Nussbaum, "Unsupervised clustering of hyperspectral paper data using t-sne," *Journal of Imaging*, vol. 6, no. 5, 2020. DOI: 10.3390/jimaging6050029 (cit. on p. 8).

[44] E. Pouyet, N. Rohani, A. K. Katsaggelos, O. Cossairt, and M. Walton, "Innovative data reduction and visualization strategy for hyperspectral imaging datasets using t-sne approach," *Pure and Applied Chemistry*, vol. 90, no. 3, pp. 493–506, 2018. DOI: 10.1515/pac-2017-0907 (cit. on p. 8).

[45] L. Van Der Maaten, "Barnes-hut-sne," *arXiv preprint arXiv:1301.3342*, 2013. DOI: 10.48550/arXiv.1301.3342 (cit. on p. 8).

[46] A. A. Goodman, "Principles of high-dimensional data visualization in astronomy," *Astronomische Nachrichten*, vol. 333, no. 5-6, pp. 505–514, 2012. DOI: 10.48550/arXiv.1205.4747 (cit. on p. 9).

[47] M. Brehmer, M. Sedlmair, S. Ingram, and T. Munzner, "Visualizing dimensionally-reduced data: Interviews with analysts and a characterization of task sequences," in *Proceedings of ACM BELIV Workshop*, 2014, pp. 1–8. DOI: 10.1145/2669557.2669559 (cit. on p. 9).

[48] M. Sedlmair, M. Brehmer, S. Ingram, and T. Munzner, “Dimensionality reduction in the wild : Gaps and guidance,” 2012 (cit. on p. 9).

[49] D. Sacha *et al.*, “Visual interaction with dimensionality reduction: A structured literature analysis,” *IEEE transactions on visualization and computer graphics*, vol. 23, no. 1, pp. 241–250, 2016. DOI: 10.1109/TVCG.2016.2598495 (cit. on p. 9).

[50] T. Höllt *et al.*, “Cytosplore: Interactive immune cell phenotyping for large single-cell datasets,” in *Computer Graphics Forum*, vol. 35, 2016, pp. 171–180. DOI: 10.1111/cgf.12893 (cit. on pp. 9, 22).

[51] A. Somarakis, V. Van Unen, F. Koning, B. Lelieveldt, and T. Höllt, “Imacyte: Visual exploration of cellular micro-environments for imaging mass cytometry data,” *IEEE transactions on visualization and computer graphics*, vol. 27, no. 1, pp. 98–110, 2019. DOI: 10.1109/TVCG.2019.2931299 (cit. on p. 9).

[52] D. Schapiro *et al.*, “Histocat: Analysis of cell phenotypes and interactions in multiplex image cytometry data,” *Nature methods*, vol. 14, no. 9, pp. 873–876, 2017. DOI: 10.1038/nmeth.4391 (cit. on p. 9).

[53] S. Warchol *et al.*, “Visinity: Visual spatial neighborhood analysis for multiplexed tissue imaging data,” *IEEE transactions on visualization and computer graphics*, 1912. DOI: 10.1109/TVCG.2022.3209378 (cit. on p. 9).

[54] R. Krueger *et al.*, “Facetto: Combining unsupervised and supervised learning for hierarchical phenotype analysis in multi-channel image data,” *IEEE transactions on visualization and computer graphics*, vol. 26, no. 1, pp. 227–237, 2019. DOI: 10.1109/TVCG.2019.2934547 (cit. on p. 9).

[55] J. Jessup *et al.*, “Scope2screen: Focus+ context techniques for pathology tumor assessment in multivariate image data,” *IEEE transactions on visualization and computer graphics*, vol. 28, no. 1, pp. 259–269, 2021. DOI: 10.1109/TVCG.2021.3114786 (cit. on p. 9).

[56] A. Somarakis *et al.*, “Visual cohort comparison for spatial single-cell omics-data,” *arXiv preprint arXiv:2006.05175*, 2020. DOI: 10.1109/TVCG.2020.3030336 (cit. on p. 9).

[57] M. Meyer, T. Munzner, A. DePace, and H. Pfister, “Multeesum: A tool for comparative spatial and temporal gene expression data,” *IEEE transactions on visualization and computer graphics*, vol. 16, no. 6, pp. 908–917, 2010. DOI: 10.1109/TVCG.2010.137 (cit. on p. 9).

[58] R. A. Leite *et al.*, “Eva: Visual analytics to identify fraudulent events,” *IEEE transactions on visualization and computer graphics*, vol. 24, no. 1, pp. 330–339, 2017. DOI: 10.1109/TVCG.2017.2744758 (cit. on p. 9).

[59] A. Krekeler, “A study of pieter de hooch’s painting technique,” in *Pieter de Hooch in Delft. From the Shadow of Vermeer*, 2019, pp. 56–79 (cit. on pp. 11, 12, 19).

[60] J. Striova, A. Dal Fovo, and R. Fontana, “Reflectance imaging spectroscopy in heritage science,” *La Rivista del Nuovo Cimento*, vol. 43, no. 10, pp. 515–566, 2020. DOI: 10.1007/s40766-020-00011-6 (cit. on p. 12).

[61] M. Alfeld, “Ma-xrf for historical paintings: State of the art and perspective,” *Microscopy and Microanalysis*, vol. 26, no. S2, pp. 72–75, 2020. DOI: 10.1017/S1431927620013288 (cit. on p. 12).

[62] M. Brehmer and T. Munzner, “A multi-level typology of abstract visualization tasks,” *IEEE transactions on visualization and computer graphics*, vol. 19, no. 12, pp. 2376–2385, 2013. DOI: 10.1109/TVCG.2013.124 (cit. on p. 14).

[63] D. Comaniciu and P. Meer, “Mean shift analysis and applications,” in *Proceedings of the seventh IEEE international conference on computer vision*, IEEE, vol. 2, 1999, pp. 1197–1203. DOI: 10.1109/ICCV.1999.790416 (cit. on p. 22).

[64] O. A. De Carvalho and P. R. Meneses, “Spectral correlation mapper (scm): An improvement on the spectral angle mapper (sam),” in *Summaries of the 9th JPL Airborne Earth Science Workshop, JPL Publication 00-18*, vol. 9, 2000 (cit. on pp. 23, 31).

[65] *Qt*, https://www.qt.io/, Accessed: 2022-06-23 (cit. on p. 30).

[66] M. Woo, J. Neider, T. Davis, and D. Shreiner, *OpenGL programming guide: the official guide to learning OpenGL, version 1.2*. Addison-Wesley Longman Publishing Co., Inc., 1999 (cit. on p. 30).

[67] N. Pezzotti *et al.*, "Gpgpu linear complexity t-sne optimization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 1, pp. 1172–1181, 2020. DOI: 10.1109/TVCG.2019.2934307 (cit. on p. 30).

[68] M. Bostock, V. Ogievetsky, and J. Heer, "D^3 data-driven documents," *IEEE transactions on visualization and computer graphics*, vol. 17, no. 12, pp. 2301–2309, 2011. DOI: 10.1109/TVCG.2011.185 (cit. on p. 30).

[69] *Qt webengine*, https://doc.qt.io/qt-6/qtwebengine-index.html, Accessed: 2022-06-23 (cit. on p. 30).

[70] M. Vijaymeena and K. Kavitha, "A survey on similarity measures in text mining," *Machine Learning and Applications: An International Journal*, vol. 3, no. 2, pp. 19–28, 2016. DOI: 10.5121/mlaij.2016.3103 (cit. on p. 32).

[71] S. Fletcher, M. Z. Islam, *et al.*, "Comparing sets of patterns with the jaccard index," *Australasian Journal of Information Systems*, vol. 22, 2018. DOI: 10.3127/ajis.v22i0.1538 (cit. on p. 32).

[72] R. R. Shamir, Y. Duchin, J. Kim, G. Sapiro, and N. Harel, "Continuous dice coefficient: A method for evaluating probabilistic segmentations," *arXiv preprint arXiv:1906.11031*, 2019. DOI: 10.48550/arXiv.1906.11031 (cit. on p. 32).

[73] I. S. Dhillon, "Co-clustering documents and words using bipartite spectral graph partitioning," in *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, 2001, pp. 269–274. DOI: 10.1145/502512.502550 (cit. on p. 33).

[74] *K-means clustering*, https://github.com/harthur/clustering, Accessed: 2022-06-23 (cit. on p. 33).

[75] *Microsoft forms*, https://forms.office.com/, Accessed: 2022-06-23 (cit. on p. 35).

[76] A. Joshi, S. Kale, S. Chandel, and D. K. Pal, "Likert scale: Explored and explained," *British journal of applied science & technology*, vol. 7, no. 4, p. 396, 2015. DOI: 10.9734/BJAST/2015/14975 (cit. on pp. 35, 41).

[77] J. Brooke, "Sus: A quick and dirty usability scale," *Usability Eval. Ind.*, vol. 189, 1995 (cit. on pp. 35, 41).

[78] T. Kleynhans, D. W. Messinger, and J. K. Delaney, "Towards automatic classification of diffuse reflectance image cubes from paintings collected with hyperspectral cameras," *Microchemical Journal*, vol. 157, p. 104 934, 2020. DOI: 10.1016/j.microc.2020.104934 (cit. on pp. 49, 50).

[79] Y. Wang, L. Chen, J. Jo, and Y. Wang, "Joint t-sne for comparable projections of multiple high-dimensional datasets," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 1, pp. 623–632, 2022. DOI: 10.1109/TVCG.2021.3114765 (cit. on p. 50).