

Visual Analysis of RIS Data for Endmember Selection

Popa, A.; Gabrieli, F.; Kroes, T.; Krekeler, A.; Alfeld, M.; Lelieveldt, B.; Eisemann, E.; Höllt, T.

DOI

[10.2312/gch.20221233](https://doi.org/10.2312/gch.20221233)

Publication date

2022

Document Version

Final published version

Published in

Proceedings Eurographics Workshop on Graphics and Cultural Heritage

Citation (APA)

Popa, A., Gabrieli, F., Kroes, T., Krekeler, A., Alfeld, M., Lelieveldt, B., Eisemann, E., & Höllt, T. (2022). Visual Analysis of RIS Data for Endmember Selection. In R. Pintus, & F. Ponchio (Eds.), *Proceedings Eurographics Workshop on Graphics and Cultural Heritage* (pp. 103-106). The Eurographics Association. <https://doi.org/10.2312/gch.20221233>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Visual Analysis of RIS Data for Endmember Selection

A. Popa¹ , F. Gabrieli² , T. Kroes³ , A. Krekeler², M. Alfeld¹ , B. Lelieveldt³ , E. Eisemann¹ , and T. Höllt¹ 

¹Delft University of Technology, Delft, The Netherlands

²Conservation and Science Department, Rijksmuseum, Amsterdam, The Netherlands

³Leiden University Medical Center, Leiden, The Netherlands

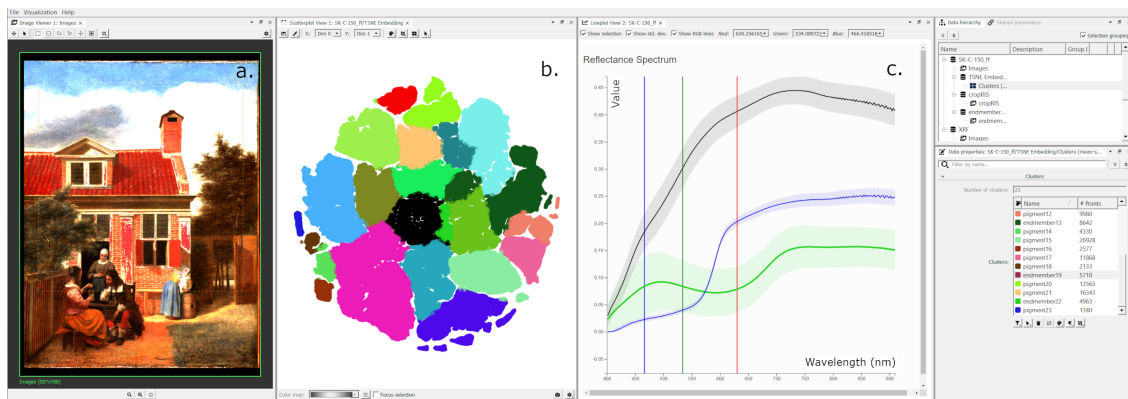


Figure 1: Overview of the implemented system showing the painting *Figures in a Courtyard behind a House by Pieter De Hooch [Kre19]*; The pixel selection is highlighted in red in the a) image view, b) embedding view and shown as the top spectrum in the c) spectrum view.

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 [BPPO5]. 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 end-member 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 $0.35 \times 0.35 \text{ mm}^2$. 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, indicated 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

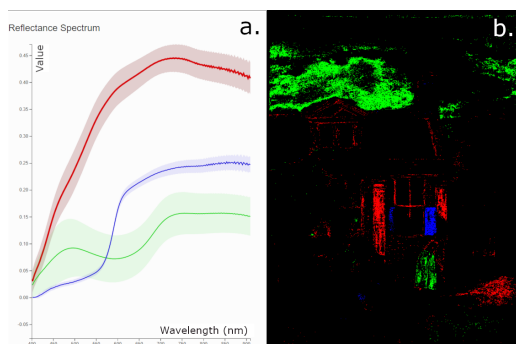


Figure 2: Endmember maps corresponding to three endmembers.

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.

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

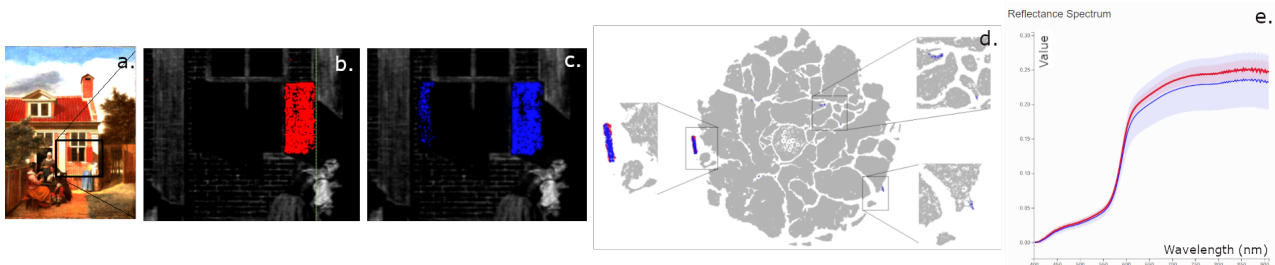


Figure 3: Comparison of t-SNE cluster (red) and endmember-based (blue) map.

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. “D³ 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, MASARCYK, 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, THIJSSSEN, 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.